

SENSORS & TRANSDUCERS

ISSN 1726-5479

vol. 272

1 / 26



Sensors & Actuators for Drones and UAVs

International Frequency Sensor Association Publishing



Sensors & Transducers

**International Official Open Access Journal of the
International Frequency Sensor Association (IFSA)
Devoted to Research and Development
of Sensors and Transducers**

Volume 272, Issue 1, April 2026

Editor-in-Chief

Prof., Dr. Sergey Y. YURISH



IFSA Publishing: Barcelona • Toronto

Sensors & Transducers is an open access journal which means that all content (article by article) is freely available without charge to the user or his/her institution. Users are allowed to read, download, copy, distribute, print, search, or link to the full texts of the articles, or use them for any other lawful purpose, without asking prior permission from the publisher or the author. This is in accordance with the BOAI definition of open access. Authors who publish articles in *Sensors & Transducers* journal retain the copyrights of their articles. The *Sensors & Transducers* journal operates under the Creative Commons License CC-BY.

Notice: No responsibility is assumed by the Publisher for any injury and/or damage to persons or property as a matter of products liability, negligence or otherwise, or from any use or operation of any methods, products, instructions or ideas contained in the material herein.

Published by International Frequency Sensor Association (IFSA) Publishing. Printed in the USA.



Contents

Volume 272
Issue 1
April 2026

www.sensorsportal.com

ISSN 2306-8515
e-ISSN 1726-5479

Research Articles

Design and Experimental Validation of a Lightweight UAV Sensing System for Geo-Referenced Air-Quality Monitoring <i>M. S. Saha, S. Baccar, M. Kadi and J. F. Carpentier</i>	1
An End-to-End Autonomous UAV System for Disaster Response: Terrain-Aware Detection, Payload Delivery, and ROS-Gazebo Validation <i>Aryan Sirsavkar, Muzammil Shaikh, Anuj Kale, Abhijeet Thore and Mayur Patil</i>	13
Towards Offline Navigation for Small UAVs Using Lightweight CNNs and Satellite Maps <i>Ludwig Kempf, Mark Umansky, Kevin B. Kochersberger</i>	23
Joint Rotor Speed Scheduling and Trajectory Optimization for Energy-Efficient Electric UAV Rotorcraft Missions <i>Emeka Chijioke and Marcin Żugaj</i>	36
A Transport-Layer-Oriented Servo Actuation Framework for UAV Systems Using USB-CAN Communication <i>Jinyoung Lee and Sungwook Cho</i>	50
A Security-Aware Edge Artificial Intelligence Framework for Robust Autonomous Decision-Making in Unmanned Aerial Systems <i>Sérgio Silva</i>	59
Civil-Military Knowledge Flows and the Diffusion of UAV Technologies under Economic Security Conditions <i>K. Miwa and K. Sumikura</i>	69

Authors are encouraged to submit article in MS Word (doc) and Acrobat (pdf) formats
by e-mail: editor@sensorsportal.com. Please visit journal's webpage with preparation instructions:
https://sensorsportal.com/sensors_and_transducers.html



Editors-in-Chief: Professor, Dr. Sergey Y. Yurish, tel.: +34 696067716, e-mail: editor@sensorsportal.com

Editors for Western Europe

Meijer, Gerard C.M., Delft Univ. of Technology, The Netherlands
Ferrari, Vittorio, Università di Brescia, Italy
Mescheder, Ulrich, Univ. of Applied Sciences, Furtwangen, Germany

Editor for Eastern Europe

Sachenko, Anatoly, Ternopil National Economic University, Ukraine

Editors for North America

Katz, Evgeny, Clarkson University, USA
Datskos, Panos G., Oak Ridge National Laboratory, USA
Fabien, J. Josse, Marquette University, USA

Editor for Africa

Maki K., Habib, American University in Cairo, Egypt

Editors South America

Costa-Felix, Rodrigo, Inmetro, Brazil
Walsoe de Reça, Noemi Elisabeth, CINSO-CITEDEF
UNIDEF (MINDEF-CONICET), Argentina

Editors for Asia

Ohyama, Shinji, Tokyo Institute of Technology, Japan
Zhengbing, Hu, Huazhong Univ. of Science and Technol., China
Li, Gongfa, Wuhan Univ. of Science and Technology, China

Editor for Asia-Pacific

Mukhopadhyay, Subhas, Massey University, New Zealand

International Advisory Board

Abdul Rahim, Ruzairi, Universiti Teknologi, Malaysia
Abramchuk, George, Measur. Tech. & Advanced Applications, Canada
Aluri, Geetha S., Globalfoundries, USA
Ascoli, Giorgio, George Mason University, USA
Atalay, Selcuk, Inonu University, Turkey
Atghiaee, Ahmad, University of Tehran, Iran
Augutis, Vygtantas, Kaunas University of Technology, Lithuania
Ayesh, Aladdin, De Montfort University, UK
Baliga, Shankar, B., Laser Components DG, Inc., USA
Barlingay, Ravindra, Larsen & Toubro - Technology Services, India
Basu, Sukumar, Jadavpur University, India
Bousbia-Salah, Mounir, University of Annaba, Algeria
Bouvet, Marcel, University of Burgundy, France
Campanella, Luigi, University La Sapienza, Italy
Carvalho, Vitor, Minho University, Portugal
Changhai, Ru, Harbin Engineering University, China
Chen, Wei, Hefei University of Technology, China
Cheng-Ta, Chiang, National Chia-Yi University, Taiwan
Cherstvy, Andrey, University of Potsdam, Germany
Chung, Wen-Yaw, Chung Yuan Christian University, Taiwan
Cortes, Camilo A., Universidad Nacional de Colombia, Colombia
D'Amico, Arnaldo, Università di Tor Vergata, Italy
De Stefano, Luca, Institute for Microelectronics and Microsystem, Italy
Ding, Jianning, Changzhou University, China
Djordjević, Alexander, City University of Hong Kong, Hong Kong
Donato, Nicola, University of Messina, Italy
Dong, Feng, Tianjin University, China
Erkmen, Aydan M., Middle East Technical University, Turkey
Fezari, Mohamed, Badji Mokhtar Annaba University, Algeria
Gaura, Elena, Coventry University, UK
Gole, James, Georgia Institute of Technology, USA
Gong, Hao, National University of Singapore, Singapore
Gonzalez de la Rosa, Juan Jose, University of Cadiz, Spain
Goswami, Amarjyoti, Kaziranga University, India
Guillet, Bruno, University of Caen, France
Hadjiloucas, Sillas, The University of Reading, UK
Hao, Shiyang, Michigan State University, USA
Hui, David, University of New Orleans, USA
Jaffrezic-Renault, Nicole, Claude Bernard University Lyon 1, France
Jamil, Mohammad, Qatar University, Qatar
Kaniusas, Eugenijus, Vienna University of Technology, Austria
Kim, Min Young, Kyungpook National University, Korea
Kumar, Arun, University of Delaware, USA
Lay-Ekuakille, Aime, University of Lecce, Italy
Li, Fengyuan, HARMAN International, USA
Li, Jingsong, Anhui University, China
Li, Si, GE Global Research Center, USA
Lin, Paul, Cleveland State University, USA
Liu, Aihua, Chinese Academy of Sciences, China
Liu, Chenglian, Long Yan University, China
Liu, Fei, City College of New York, USA
Mahadi, Muhammad, University Tun Hussein Onn Malaysia, Malaysia
Mansor, Muhammad Naufal, University Malaysia Perlis, Malaysia

Marquez, Alfredo, Centro de Investigacion en Materiales Avanzados, Mexico
Mishra, Vivekanand, National Institute of Technology, India
Moghavvemi, Mahmoud, University of Malaya, Malaysia
Morello, Rosario, University "Mediterranea" of Reggio Calabria, Italy
Mulla, Imtiaz Sirajuddin, National Chemical Laboratory, Pune, India
Nabok, Aleksey, Sheffield Hallam University, UK
Neshkova, Milka, Bulgarian Academy of Sciences, Bulgaria
Pal, Jitendra, Carnegie Mellon University, USA
Passaro, Vittorio M. N., Politecnico di Bari, Italy
Patil, Devidas Ramrao, R. L. College, Parola, India
Penza, Michele, ENEA, Italy
Pereira, Jose Miguel, Instituto Politecnico de Seteabal, Portugal
Pillariseti, Anand, Sensata Technologies Inc, USA
Pogacnik, Lea, University of Ljubljana, Slovenia
Pullini, Daniele, Centro Ricerche FIAT, Italy
Qiu, Liang, Avago Technologies, USA
Reig, Candid, University of Valencia, Spain
Restivo, Maria Teresa, University of Porto, Portugal
Rodríguez Martínez, Angel, Universidad Politécnic de Cataluña, Spain
Sadana, Ajit, University of Mississippi, USA
Sadeghian Marnani, Hamed, TU Delft, The Netherlands
Sapozhnikova, Ksenia, D. I. Mendeleev Institute for Metrology, Russia
Singhal, Subodh Kumar, National Physical Laboratory, India
Shah, Kriyang, La Trobe University, Australia
Shi, Wendian, California Institute of Technology, USA
Shmaliy, Yuriy, Guanajuato University, Mexico
Song, Xu, An Yang Normal University, China
Srivastava, Arvind K., Systron Donner Inertial, USA
Stefanescu, Dan Mihai, Romanian Measurement Society, Romania
Sumridetchkajorn, Sarun, Nat. Electr. & Comp. Tech. Center, Thailand
Sun, Zhiqiang, Central South University, China
Sysoev, Victor, Saratov State Technical University, Russia
Thirunavukkarasu, I., Manipal University Karnataka, India
Thomas, Sadiq, Heriot Watt University, Edinburgh, UK
Tian, Lei, Xidian University, China
Tianxing, Chu, Research Center for Surveying & Mapping, Beijing, China
Vanga, Kumar L., ePack, Inc., USA
Vazquez, Carmen, Universidad Carlos III Madrid, Spain
Wang, Jiangping, Xian Shiyong University, China
Wang, Peng, Qualcomm Technologies, USA
Wang, Zongbo, University of Kansas, USA
Xu, Han, Measurement Specialties, Inc., USA
Xu, Weihe, Brookhaven National Lab, USA
Xue, Ning, Agiltron, Inc., USA
Yang, Dongfang, National Research Council, Canada
Yang, Shuang-Hua, Loughborough University, UK
Yaping Dan, Harvard University, USA
Yue, Xiao-Guang, Shanxi University of Chinese Traditional Medicine, China
Xiao-Guang, Yue, Wuhan University of Technology, China
Zakaria, Zulkarnay, University Malaysia Perlis, Malaysia
Zhang, Weiping, Shanghai Jiao Tong University, China
Zhang, Wenming, Shanghai Jiao Tong University, China
Zhang, Yudong, Nanjing Normal University China

12th International Conference on Sensors and Electronic Instrumentation Advances

<https://www.seia-conference.com/>

30 September - 2 October 2026
Granada, Spain



The XII International Conference on Sensors and Electronic Instrumentation Advances brings together researchers, innovators, and industry leaders to advance sensor science and electronic instrumentation by linking fundamental research, engineering practice, and real-world deployment. The conference provides a platform for presenting innovations that improve the measurement, monitoring, and control of complex systems in industry, environment, healthcare, mobility, energy, and smart infrastructure. Particular emphasis is placed on moving sensing technologies from prototypes to scalable, reliable applications. SEIA' 2026 welcomes original research papers, applied case studies, and technology demonstrations across diverse sensor types and modalities. Main topics include accuracy, calibration, uncertainty evaluation, reliability, validation in demanding environments, manufacturability, testing, interoperability, system integration, secure connectivity, and regulatory compliance. Through technical sessions and applied tutorials, the conference aims to strengthen the global sensing community and promote secure, interoperable, and standards-aware solutions with practical impact.

Sensor Technology

- Accelerometers
- Inclometers
- Gyroscopes
- Mechanical Sensors
- Optical Sensors
- Optical Fiber Sensors
- Photonic Sensors
- Chemical Sensors
- Biosensors
- Immunosensors
- BioMEMS
- Temperature Sensors
- Pressure Sensors
- Acoustic Sensors
- Electromagnetic Sensors
- Gas Sensors
- Humidity Sensors
- Infrared Sensors, Devices and Thermography
- Radiation Sensors
- Multi Sensor Fusion
- Smart Sensors
- Intelligent Sensors
- Artificial Intelligence based Sensors and Sensor Systems
- Virtual Sensors
- Sensor Interfacing and Signal Conditioning
- Sensor Calibration
- Nanomaterials and Electronics Technology for Sensors
- Semiconductor Materials for Sensors
- Polymer Materials for Sensors
- MEMS and NEMS
- Remote Sensors and Telemetry
- Sensor Applications

Sensor Instrumentation and Measuring Technology

- Metrology and Measurement Science
- Methods of Measurements
- Calibrations and Standards
- Measurement of Electrical Quantities
- Time and Frequency Measurements
- Measurement of Force, Mass, Torque, Inclination and Acceleration
- Magnetic Measurements
- Hardness Measurements
- Measurement of Geometrical Quantities
- Temperature and Thermal Measurements
- Pressure and Vacuum Measurements
- Vibration and Noise Measurement
- Flow Measurements
- Chemical Measurements
- Quantum Measurements and Photonics
- Acoustics and the Ultrasonic Measurements
- Environmental Measurements
- Power and Energy Measurements
- Measurement of Human Functions
- Measurements in Biology and Medicine
- Mathematical Tools for Measurements
- Optical and Radiation Measurements
- Microwave Measurements
- Virtual Instruments and Data Acquisition Systems
- Soft Measurements
- Measurement Systems
- Distributed Measurements
- Analog-to-Digital Converters, Digital and Mixed Signal Processing
- Waveform Analysis and Measurements
- Scientific and Industrial Instrumentation
- Cyber-Physical Systems and IoT
- Experimental Mechanics

- Measurement in Robotics
- Metrology in Food and Nutrition
- Intelligent and Computer Vision Instruments
- Reliability of Instrument and Measurement Systems
- Nanometrology
- Technical Diagnostics and Testing
- Nondestructive Testing
- Blockchain Applications in Metrology and Measurements
- Education and Training in Measurement and Instrumentation

Special Sessions:

Authors are welcome to propose and manage special sessions during the SEIA' 2026. Each special session will contain 4-6 papers in a related field as specified above.

Session organizers will get: Certificate of Appreciation; Free Registration; Special Publishing Theme within Conference Proceedings.

Contribution Types:

- Keynote Presentations
- Invited Talks
- Special Sessions
- Regular Sessions
- Poster Sessions
- Special Sessions
- Online Sessions
- Exhibition

Deadlines:

Submission (2-page extended abstract): **1 June 2026**
Notification of acceptance: **30 June 2026**
Registration: **25 July 2026**
Camera ready and late registration: **30 August 2026**



12th International Conference on Sensors and Electronic Instrumentation Advances



<https://www.seia-conference.com/>

One event - three different publications !

1) All registered abstracts will be published in the conference proceedings (with the ISBN and DOI). Some previous SEIA Conference Series Proceedings are indexed in Conference Proceedings Citation Index (CPCI), Web of Science by Clarivate Analytics.

2) Authors will be invited to submit full-page extended papers to special issue of open access journal: *Sensors & Transducers* (ISSN: 2306-8515, e-ISSN: 1726-5479), published by IFSA Publishing (Barcelona, Spain), and indexed in Scopus and MDPI *Sensors* (ISSN 1424-8220) open access journal, published by MDPI AG (Basel, Switzerland). Impact factor: 3.5.

3) The papers, presented at SEIA' 2026 may be extended to open access book chapters for the Book Series on 'Advances in Sensors', Vol. 10, 'Advances in Biosensors', Vol. 4 or 'Advances in Measurements and Instrumentation', Vol. 3. Some previous Book Series volumes are indexed in the Book Citation Index (Wed of Science) by Clarivate Analytics.

Conference Hotel

The SEIA' 2026 Conference will be held in Barceló Granada Congress Hotel, Granada, Spain

Address: Calle Maestro Montero 12 18004 Granada Spain

Social Programme

- Welcome Cocktail: 29 September 2026 (20:00-21:30)
- Gala Dinner: 1 October 2026 (20:00-23:00)
- Farewell Cocktail: 2 October 2026, Friday (after lunch time)

Organizing Committee

Chairman

Prof., Dr. Sergey Y. Yurish (IFSA, Spain),
e-mail: seia@sensorsportal.com

Advisory Chairman

Prof. Vijyakumar Varadarajan
(Hano Technical University, Somalia)

Conference and Publication Manager

Tetyana Zakharchenko (IFSA Publishing, S.L., Spain)

Steering Committee:

Prof. Gennaro Conte (University Roma Tre, Italy)

Dr. Pavel Shuk (Saint-Gobain NorPro, USA)

Dr. Marius Gheorghe (Ideal Aerosmith, Inc., USA)

Dr. Paolo Dabove (Politecnico di Torino, Italy)

Prof. George Kiriakidis (European Materials Research Society (E-MRS), France)

Prof., Dr. Arkady Zhukov (Univ. of the Basque Country, Spain)



Design and Experimental Validation of a Lightweight UAV Sensing System for Geo-Referenced Air-Quality Monitoring

¹ M. S. SAHA, ¹ S. BACCAR, ¹ M. KADI and ² J. F. CARPENTIER

¹ IRSEEM, ESIGELEC, University of Rouen Normandy, Rouen, France

² ELISA Aerospace, Sain-Quentin, France

E-mail: Moni-Sankar.Saha@esigelec.fr, Sahbi.Baccar@esigelec.fr, Moncef.Kadi@esigelec.fr, j.carpentier@elisa-aerospace.fr

Received: 19 Feb. 2026 /Revised: 13 April 2026 /Accepted: 21 April 2026 /Published: 28 April 2026

Abstract: Unmanned Aerial Vehicles (UAVs) equipped with low-cost sensing platforms are increasingly employed for environmental monitoring due to their mobility, rapid deployment capability, and ability to acquire spatially distributed measurements in areas that are difficult to access using conventional ground-based instrumentation. This article presents a lightweight multi-sensor UAV platform based on the ESP32 microcontroller, designed for low-cost and flexible environmental sensing applications. The proposed system integrates four sensing modules: an MPU6050 inertial measurement unit (IMU) for attitude estimation, a NEO-6M GPS receiver for geolocation, an MQ135 gas sensor for air-quality monitoring, and an ESP32-CAM module for visual documentation. A unified data acquisition and synchronization framework is developed to align heterogeneous sensor streams operating at different sampling rates. Sensor-specific compensation strategies are implemented to mitigate IMU bias and drift, gas-sensor stabilization effects, and GPS measurement noise. In addition, measurement uncertainty is quantified following the principles of the Guide to the Expression of Uncertainty in Measurement (GUM), providing a systematic evaluation of the reliability of the sensed data. The proposed system is experimentally validated through UAV flight tests, demonstrating stable operation and consistent multi-sensor integration. The results show coherent reconstruction of UAV trajectory and attitude, as well as repeatable variations in gas sensor response correlated with UAV motion and spatial position. These findings demonstrate the feasibility of the proposed sensor synchronization and fusion framework for low-cost UAV-based environmental sensing, although comparative evaluation against alternative fusion baselines is left for future work.

Keywords: UAV instrumentation, ESP32, Sensor fusion, IMU calibration, Gas sensing, GPS tracking, Environmental monitoring.

1. Introduction

Unmanned Aerial Vehicles (UAVs) have emerged as a powerful platform for environmental monitoring due to their mobility, flexibility, and ability to provide spatially distributed measurements. Unlike fixed ground-based stations, UAVs can access hard-to-reach or hazardous environments and follow programmable trajectories, enabling the collection of geo-referenced data with high spatial resolution. This capability makes UAVs particularly suitable for air-quality

monitoring, atmospheric sensing, infrastructure inspection, and precision agriculture applications [1].

Recent advances in embedded systems and low-power electronics have facilitated the integration of multiple sensors into lightweight UAV platforms. Low-cost microcontrollers such as the ESP32 enable real-time acquisition of heterogeneous data from inertial sensors, gas sensors, and positioning systems, making it possible to design compact and energy-efficient airborne sensing payloads. Consequently, multi-sensor UAV systems have been

increasingly investigated for environmental monitoring, where spatial and temporal variations of atmospheric parameters must be captured simultaneously.

Despite these advancements, several technical challenges remain in achieving reliable and accurate measurements using low-cost UAV-based sensing systems. Inertial Measurement Units (IMUs) suffer from drift and bias accumulation, which degrade orientation estimation over time. Gas sensors, particularly metal-oxide semiconductor (MOS) sensors such as the MQ135, exhibit nonlinear responses, temperature dependence, and warm-up effects that affect measurement stability. Global Navigation Satellite System (GPS) receivers provide absolute positioning but are subject to noise, multipath effects, and limited update rates. Furthermore, vibrations induced by UAV motors and propellers introduce additional disturbances in sensor readings. Another critical issue is the asynchronous nature of data acquisition, as different sensors operate at different sampling frequencies, leading to temporal misalignment and inconsistency in multi-sensor datasets [2].

To address these limitations, recent work has explored sensor fusion techniques that combine IMU, GPS, and environmental sensors to improve measurement accuracy and robustness. Data synchronization strategies, filtering methods, and compensation models have been proposed to mitigate the impact of noise, drift, and environmental variability. However, many existing solutions rely on high-cost instrumentation or computationally intensive algorithms, limiting their applicability in lightweight and low-cost UAV platforms. In addition, comprehensive experimental validation combining spatial mapping, temporal synchronization, and uncertainty analysis remains limited in the literature.

In this work, a lightweight and low-cost multi-sensor UAV sensing framework is proposed, based on an ESP32 embedded platform. The system integrates four sensing modules, including an IMU for attitude estimation, a gas sensor for air-quality monitoring, a GPS receiver for geo-referencing, and a vision module for event detection, as illustrated in Fig. 1. A unified acquisition and synchronization framework is developed to align heterogeneous sensor data that operate at different sampling rates. In addition, compensation and uncertainty-aware processing methods are applied to mitigate the effects of sensor bias, drift, and environmental disturbances.

The proposed system is experimentally validated through UAV flight tests, demonstrating stable operation and the capability to generate coherent spatiotemporal environmental measurements. The results demonstrate the feasibility of the proposed low-complexity sensor-fusion approach for synchronized multi-sensor environmental sensing on a low-cost UAV platform. [3].

The remainder of this article is organized as follows. Section 2 describes individual sensor characterization. Section 3 presents the overall

architecture of the UAV system. Section 4 introduces the data synchronization and sensor-fusion framework. Finally, Section 5 concludes the article.




Fig. 1. Overview of the UAV Environmental Monitoring System with Onboard IMU, MQ135 Gas Sensor, NEO-6M GPS, and ESP32-CAM.

A preliminary version of this work was previously presented in the DAUS 2026 conference proceedings and is cited in this manuscript as [16]. The present article is a substantially extended journal version. Compared with the earlier conference paper, this manuscript provides a broader description of the complete ESP32-based UAV sensing platform, including the MPU6050 IMU, Neo-6M GPS, MQ135 gas sensor, and ESP32-CAM; expanded standalone characterization of the sensing modules; a more explicit formulation of the data synchronization and sensor-fusion framework; and extended experimental validation through synchronized flight-data analysis, trajectory reconstruction, altitude-correlated gas-response analysis, and geo-referenced environmental mapping. These additions significantly expand the methodological detail, experimental discussion, and overall scope beyond the prior conference dissemination.

2. Standalone Sensor Tests and Calibration

Prior to multi-sensor fusion, each module was evaluated independently to verify functionality, quantify noise, and derive calibration parameters. This section details the test protocols, metrics, and governing equations for the IMU (MPU6050), gas sensor (MQ135), GPS (Neo-6M), and ESP32-CAM.

2.1. IMU-Based Stability Assessment

The performance of the MPU6050 inertial measurement unit (IMU) was evaluated independently using the ESP32 acquisition platform to assess its short-term stability and dynamic response. The sensor was sampled at 20 Hz over a duration of 50 s. During the experiment, the IMU was initially placed on a flat surface under static conditions and then manually

tilted between 10 s and 15 s to introduce controlled motion.

2.1.1. Static Stability

During the static intervals (0–10 s and 20–50 s), the accelerometer outputs remain stable, with the vertical component $a_z \approx -0.95 g$ and the horizontal components $a_x, a_y \approx 0 g$, indicating proper alignment with the gravity vector. The magnitude of the acceleration vector is given by

$$\| \mathbf{a} \| = \sqrt{a_x^2 + a_y^2 + a_z^2},$$

which remains close to 1 g , with deviations limited to approximately $\pm 0.05 g$. This confirms low measurement noise and good sensor stability under static conditions.

2.1.2. Dynamic Response

During the tilting interval (10–15 s), the accelerometer signals exhibit transient variations in a_x and a_y , while a_z decreases as a function of the tilt angle, consistent with the projection of the gravity vector onto the sensor axes. Simultaneously, the gyroscope records angular-rate peaks up to $\pm 50^\circ/s$, demonstrating sufficient sensitivity to capture short-duration rotational motion [4].

The instantaneous angular-rate magnitude is expressed as

$$\| \boldsymbol{\omega} \| = \sqrt{g_x^2 + g_y^2 + g_z^2},$$

where g_x, g_y, g_z denote the gyroscope measurements along the three axes.

To quantify gyroscope stability, the bias and drift are estimated over the static intervals using

$$b_{\omega,i} = \frac{1}{N} \sum_{k=1}^N g_i(k),$$

$$\sigma_{\omega,i} = \sqrt{\frac{1}{N-1} \sum_{k=1}^N (g_i(k) - b_{\omega,i})^2},$$

where $i \in \{x, y, z\}$ and N denotes the number of samples acquired during the static period. The measured gyroscope drift remains below $0.05^\circ/s$, indicating good short-term stability.

2.1.3. Stability Summary

Table 1 summarizes the computed metrics. The accelerometer maintains a near-constant norm, while the gyroscope exhibits low drift and rapid recovery, confirming that the IMU provides sufficient stability for UAV-based sensor fusion and environmental measurements.




Fig. 2. Standalone MPU6050 response showing tri-axial acceleration (top) and angular rate (bottom). Static intervals (0–10 s and 20–50 s) demonstrate stable behavior, while peaks observed between 10 s and 15 s correspond to controlled manual perturbations

Overall, the IMU exhibits stable bias, low noise, and reliable dynamic response, ensuring accurate orientation estimation when fused with GPS and gas-sensing data during UAV operation.

Table 1. Summary of MPU6050 Stability Metrics.

Parameter	Symbol	Value	Interpretation
Accel. magnitude	$\ \mathbf{a} \ $	1.00 $\pm 0.03 g$	gravity alignment
Accel. variance	σ_a^2	$< 0.002 g^2$	High static stability
Gyro bias	$b_{\omega,i}$	$< 0.05^\circ/s$	Negligible drift
Gyro std. dev.	$\sigma_{\omega,i}$	$0.8^\circ/s$	Low-frequency noise
Gyro peak	$\ \boldsymbol{\omega} \ _{\max}$	$50^\circ/s$	Dynamic sensitivity
Recovery time	t_r	$< 3 s$	Fast stabilization

2.2. Neo-6M GPS Receiver Evaluation

The Neo-6M GPS module was first evaluated under static open-sky conditions to assess short-term positioning accuracy and fix stability. Fig. 3 shows the latitude–longitude scatter recorded over a 10 min acquisition at 1 Hz. National Marine Electronics Association (NMEA) data were parsed in MATLAB to extract position, altitude, and Horizontal Dilution of Precision (HDOP) to avoid restriction zone.

The great-circle distance between each GPS fix (φ, λ) where φ denotes the geodetic latitude and λ the geodetic longitude, expressed in degrees. and a reference position (φ_0, λ_0) was computed using the Haversine formula [2].

$$d = 2R \arcsin \sqrt{\sin^2 \frac{\Delta\varphi}{2} + \cos\varphi_0 \cos\varphi \sin^2 \frac{\Delta\lambda}{2}},$$

with earth radius $R = 6371$ km. Position precision was quantified. To quantify the positioning performance, two standard accuracy metrics are considered: the horizontal circular error probable at 50 (CEP_{50}), which characterizes the median horizontal localization error, and the 3D root mean square error (RMS_{3D}), which accounts for the combined dispersion along the three spatial axes. These metrics are expressed as

$$CEP_{50} = 0.59\sigma_x + 0.59\sigma_y,$$

$$RMS_{3D} = \sqrt{\sigma_x^2 + \sigma_y^2 + \sigma_h^2}$$

The module achieved $CEP_{50} = 2.8$ m and $RMS_{3D} < 5$ m, with altitude varying within ± 3 m, consistent with manufacturer specifications.

2.2.1. GPS Flight Data Comparison

Fig. 3 presents two representative datasets acquired during low-altitude flights. These experiments were used to assess repeatability of the Neo-6M position and velocity estimates.

Case 1: Static and Low-Dynamic Test (200 s). During the static/slow-motion experiment, the reconstructed horizontal trajectory forms a compact spatial cluster with a dispersion of approximately 5–8 m. This dispersion is primarily attributed to GPS measurement noise, satellite geometry, and receiver jitter. The estimated altitude exhibits a gradual decrease from 115 m to 70 m, which can be associated

with the internal filtering and barometric correction mechanisms of the GPS module. The measured ground speed remains within the range of 1–3 km/h, consistent with near-static motion conditions. Satellite availability varies between 7 and 9 satellites, resulting in favorable.

Case 2: Short-Duration Displacement Test (30 s). In this experiment, the UAV was manually displaced over a distance of approximately 40–50 m. The reconstructed trajectory yields a total displacement of 44.5 m, computed using the segment-wise Haversine formulation, which shows good agreement with the expected ground-truth distance. The altitude remains relatively stable within the range of 90–100 m, indicating consistent vertical positioning performance. Compared to the static test, the trajectory exhibits reduced spatial dispersion, which can be attributed to the shorter acquisition duration and improved satellite tracking conditions during motion.

Comparative Analysis. A comparison of the two test scenarios highlights that Case 2 provides higher localization reliability, characterized by reduced trajectory dispersion and lower speed variability, with a standard deviation $\sigma_v < 0.5$ km/h. In contrast, Case 1 exhibits increased positional drift due to prolonged observation time and accumulated GPS uncertainties. These results demonstrate that the ESP32–Neo-6M subsystem achieves sufficient localization accuracy for short-range UAV operations, making it suitable for geo-referenced sensing applications such as air-quality mapping and environmental monitoring.

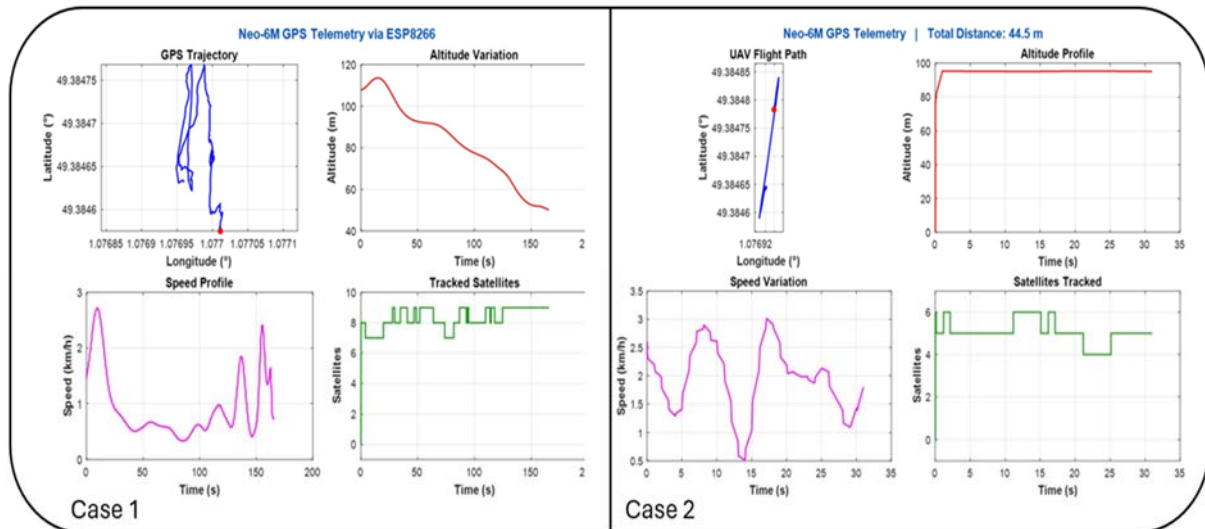


Fig. 3. Comparative GPS trajectories for Case 1 (200 s static/slow-motion) and Case 2 (30 s short displacement), showing horizontal position clusters, altitude variations, and reconstructed path length.

2.3. Auxiliary Visual Sensing Using ESP32-CAM

Besides the core IMU, GPS, and gas-sensing modules, the UAV platform also includes an

ESP32-CAM as an auxiliary visual sensor intended for contextual awareness and future multimodal extensions. This ensures robustness against motion blur, illumination variations, and sensor noise [7].

2.3.1. Color Normalization and Chromaticity Features

For each frame, the RGB pixel vector $I(x, y) = [R, G, B]$ is converted to normalized chromaticity coordinates [8]

$$r = \frac{R}{R+G+B}, g = \frac{G}{R+G+B}, b = \frac{B}{R+G+B},$$

where fire-colored regions typically satisfy $r > g > b$.

Two pixels with the same color but different brightness may have very different raw RGB values. To remove this effect, we normalize the RGB vector:

$$c = \frac{1}{\sqrt{R^2+G^2+B^2}} \begin{bmatrix} R \\ G \\ B \end{bmatrix}$$

2.3.2. Hybrid Fire Classification

A pixel is first evaluated by deterministic color rules:

$$R > T_R, G > T_G, B < T_B, R > G > B$$

To detect fire regions, an initial binary mask $M_1(x, y)$ is generated using empirical RGB thresholds $T_R = 0.6$, $T_G = 0.2$, and $T_B = 0.3$. To improve robustness against illumination changes, the normalized chromaticity vector $\mathbf{v} = [r, g, b]^T$ is evaluated using a Gaussian fire-likelihood model [3]:

$$P_{fire}(v) = \exp[-1/2 (v - \mu)^T \Sigma^{-1} (v - \mu)]$$

The Gaussian parameters were estimated from 12 manually annotated fire images comprising approximately 18000 fire pixels captured under varying illumination conditions. The sample mean and covariance were computed using maximum likelihood estimation. With mean and covariance

$$\mu = \begin{bmatrix} 0.60 \\ 0.32 \\ 0.08 \end{bmatrix}, \Sigma = \begin{bmatrix} 0.015 & 0.002 & 0.001 \\ 0.002 & 0.010 & 0.001 \\ 0.001 & 0.001 & 0.005 \end{bmatrix}$$

Pixels satisfying $P_{fire} > 0.35$ form the probabilistic mask $M_2(x, y)$. The final segmentation is obtained by combining both models:

$$M(x, y) = M_1(x, y) M_2(x, y)$$

The detected fire area is computed as

$$N_{fire} = \sum_{x,y} M(x, y),$$

and a fire alarm is raised when $N_{fire} \geq 4000$ for QVGA resolution.

To suppress transient false alarms, a temporal exponential filter is applied:

$$S(k) = \alpha N_{fire}(k) + (1 - \alpha)S(k - 1),$$

$$\alpha = 0.25,$$

and the normalized fire-confidence score is defined as

$$C_{fire} = \frac{S(k)}{76800},$$

where $76800 = 320 \times 240$ is the total number of pixels in a QVGA frame.

2.3.3. Experimental Summary

In the sequence shown in Fig. 4, the system detected $N_{fire} = 4228$ fire pixels with a confidence level of $C_{fire} = 0.055$, exceeding the alarm threshold. The results confirm that the ESP32-CAM, combined with the hybrid detection algorithm, reliably identifies fire events despite motion, vibration, and low-cost camera limitations.



Fig. 4. Fire-region segmentation mask from the proposed hybrid model.

2.4. Summary and Pass/Fail Criteria

Each sensor is evaluated according to the *Guide to the Expression of Uncertainty in Measurement* (GUM) framework if: (i) the IMU maintains $\| \mathbf{a} \| \approx 1g$ and low gyro bias, (ii) the MQ135 produces consistent R_s/R_0 variation, (iii) the GPS maintains consumer-grade precision (CEP₅₀ and Haversine error within limits), and (iv) the ESP32-CAM maintains stable timing and adequate image sharpness S . These metrics ensure reliable fusion in Section 4.

2.5. MQ135 Gas Sensor Calibration

The MQ135 gas sensor provides an analog output that varies according to the resistance change of its sensing element when exposed to different gaseous environments. In this work, the sensor response was characterized using the normalized resistance ratio R_s/R_0 , where R_s is the sensor resistance under the measured condition and R_0 is the reference resistance in clean air. The sensor resistance is obtained from the voltage-divider configuration as

$$SR_S = R_L \left(\frac{V_c - V_{out}}{V_{out}} \right),$$

where R_L is the load resistance, V_c is the supply voltage, and V_{out} is the measured sensor output voltage.

For response characterization, an empirical power-law model was used:

$$C = A \left(\frac{R_S}{R_0} \right)^B,$$

with fitted constants $A \approx 123$ and $B \approx -2.36$. This relation provides an approximate concentration scale for interpreting the sensor response. However, in the present study, the MQ135 measurements are primarily interpreted as a relative gas-response indicator rather than as a rigorously validated absolute gas-concentration measurement, since calibration against certified reference gases under controlled environmental conditions was not performed.

Fig. 5 illustrates the calibration curve used to relate the normalized sensor response R_S/R_0 to the estimated gas concentration. The decreasing trend of the curve reflects the typical behavior of metal-oxide semiconductor gas sensors, where the sensor resistance decreases as the gas concentration increases.




Fig. 5. Calibration characteristic of the MQ135 gas sensor showing the relationship between the normalized resistance ratio R_S/R_0 and an approximate concentration scale used for relative response interpretation.

2.6. Spatial Air-Quality Mapping

To demonstrate the capability of the proposed system for environmental monitoring, the UAV sensing payload was used to collect gas concentration data along a predefined flight trajectory. Each measurement was associated with the UAV position, allowing the construction of a spatial air-quality map.

Fig. 6 presents the resulting air-quality heatmap obtained from the UAV measurements. The color scale represents the corrected gas-sensing response recorded at different sampling locations. The map

highlights spatial variations in the measured gas levels and demonstrates the feasibility of the proposed low-cost UAV-based sensing platform for environmental monitoring applications.




Fig. 6. Spatial relative air-quality map generated from UAV measurements. The color scale represents the MQ135 gas-sensor response obtained at different positions along the UAV trajectory.

2.7. MQ135 Response and Calibration Analysis

Fig. 7 illustrates the temporal response of the MQ135 gas sensor interfaced with the ESP32. The output voltage V_{out} was recorded during the transitions from clean air to gas exposure and back to baseline; gray points represent raw ADC samples, and the blue curve shows filtered data filtered on a moving-average [5].

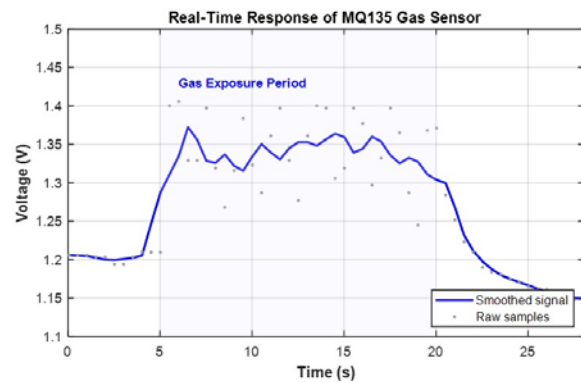


Fig. 7. Real-time MQ135 sensor response showing gas exposure and recovery. Smoothed curve overlaid on raw ADC samples.

The MQ135 operates as a variable resistor R_S in a voltage-divider configuration, with output voltage

$$V_{out} = V_{CC} \cdot \frac{R_L}{R_S + R_L},$$

where $V_{CC} = 3.3$ V. Gas exposure (5–20 s) reduces R_S due to adsorption on the CO_2 layer, increasing V_{out} , while desorption after removal causes voltage recovery.

The dynamic behavior is modeled as

$$V(t) = V_0 + \Delta V(1 - e^{-t/\tau_r}),$$

during the response phase, and

$$V(t) = V_s e^{-t/\tau_d} + V_{\infty},$$

during recovery, with time constants $\tau_r \approx 2.5$ s and $\tau_d \approx 6.0$ s.

The steady-state voltage increased from $V_0 \approx 1.20$ V to $V_s \approx 1.38$ V, yielding

$$\Delta V_{rel} = \frac{V_s - V_0}{V_0} \times 100,$$

which confirms the high sensitivity and suitability of the ESP32–MQ135 module for UAV-based environmental sensing.

3. System Design and Architecture

The proposed UAV-based measurement platform integrates an ESP32 microcontroller with multiple sensing modules for geo-referenced environmental monitoring. The system comprises [9] three main units: (i) the UAV carrier, (ii) the ESP32-based multi-sensor payload incorporating the Neo-6M GPS, MPU6050 IMU, MQ135 gas sensor, and ESP32-CAM, and (iii) the ground station for data logging and supervisory control. The hardware prototype is shown in Fig. 8.




Fig. 8. Prototype of the ESP32-based multi-sensor UAV module integrating GPS, IMU, gas sensor, and camera.

3.1. UAV Platform

MAVLink (Micro Air Vehicle Link) is a lightweight telemetry and command protocol used in UAV systems to exchange flight data between the autopilot and external devices. Here, it is used to retrieve real-time variables such as attitude, altitude, and velocity for subsequent geo-referencing and motion compensation.

3.2. Correlation of Gas Sensor Output with Altitude

The MQ135 response was analyzed as a function of UAV altitude. As shown in Fig. 9, the sensor voltage (sampled at 10 Hz) increases between 10–12 m, consistent with the imposed altitude-dependent gas and temperature gradients. The observed dispersion ($\approx \pm 50$ mV) is mainly due to ADC noise and short-term sensor fluctuations, remaining acceptable for low-cost sensing [10].

These results confirm that the ESP32-based acquisition chain provides stable and monotonic measurements correlated with flight telemetry, enabling altitude-aware air-quality profiling.

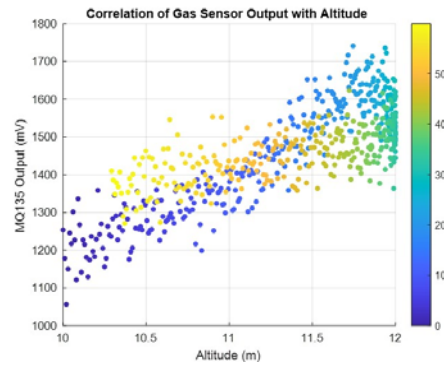


Fig. 9. Correlation between MQ135 output and altitude during flight.

3.3. Geo-Referenced Air-Quality Mapping

Geo-referenced MQ135 measurements were obtained by fusing the sensor output with latitude–longitude data from the NEO-6M GPS. Fig. 10 presents the resulting 2D heatmap, where each 10 Hz sample is color-coded according to its MQ135 voltage level, representing relative gas-response variations along the UAV path [5].

The measured flight traces a circular trajectory of approximately 60 m diameter around (49.440° N, 1.090° E). Distinct spatial variations appear along the loop, with higher voltage values (yellow/red) indicating localized increases in sensor response or turbulence-induced effects, and lower values (blue) corresponding to comparatively cleaner regions. [11].

The coherent spatial pattern further confirms correct synchronization between the ESP32 sensor stream and the GPS telemetry, validating the suitability of low-cost UAV systems for geo-referenced environmental monitoring.

3.4. System Stability and IMU Analysis

Figs. 11 and 12 show the time-domain response of the MPU6050 IMU during a 60 s measured flight,

sampled at 10 Hz and covering both steady and maneuvering intervals.

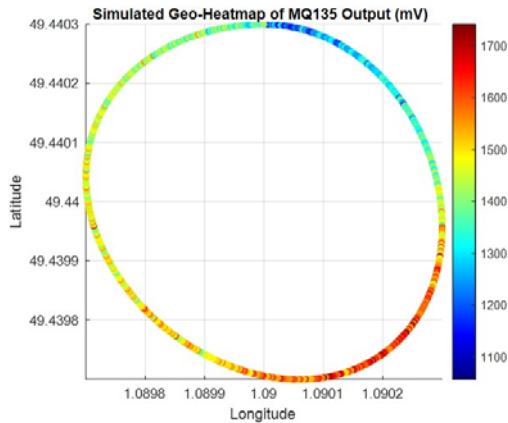


Fig. 10. Geo-referenced MQ135 measurements during circular UAV flight.

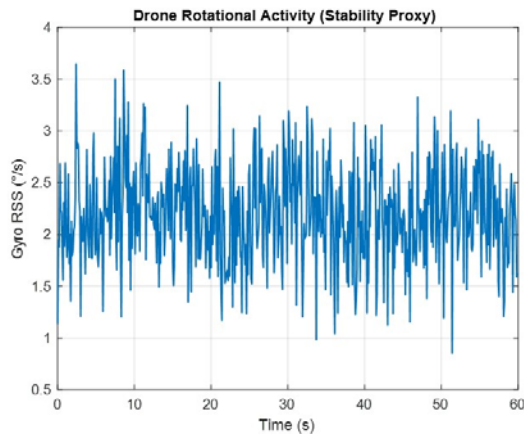


Fig. 11. MPU6050 gyroscope RSS showing 1–4°/s bounded variations and stable attitude behavior.

1) Gyroscope Response: The root-sum-square (RSS) angular velocity in Fig. 11 remains around 2.2°/s, with peaks below 4°/s, reflecting minor attitude corrections during circular flight. The absence of large transients indicates good rotational stability and limited vibration coupling to the sensing modules.

2) Accelerometer Response: Fig. 12 shows the acceleration magnitude close to 1 g, with deviations of about ± 0.05 g due to vertical motion and vibration, confirming proper IMU calibration and stable translation.

Overall, the IMU data confirm stable flight conditions suitable for reliable environmental measurements.

4. Sensor Fusion and Data Synchronization

The fusion framework in this section builds upon the independent sensor evaluations already performed

in Section 2I, where the MPU6050, MQ135, NEO-6M, and ESP32-CAM were each tested separately. These preliminary tests provided baseline calibration parameters, noise characteristics, and timing behavior for every module. As a result, the fusion and synchronization pipeline introduced here relies directly on the validated sensor outputs from Section 2, integrating them into a unified spatiotemporal representation suitable for UAV-based environmental measurement [12].

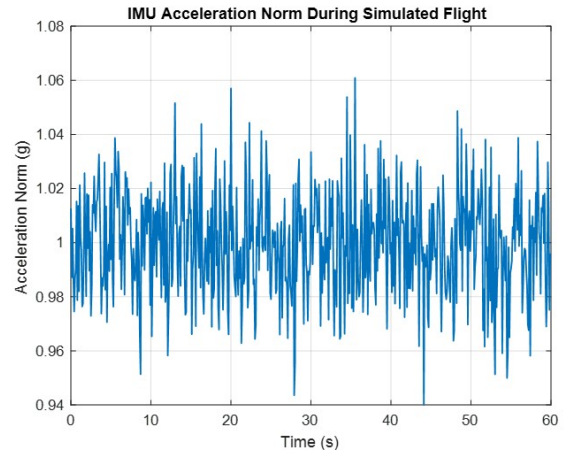


Fig. 12. MPU6050 acceleration norm staying near 1 g, indicating stable platform behavior.

4.1. Overview of the Fusion Algorithm

Fig. 13 presents the proposed sensor-fusion and synchronization pipeline. The process begins with the independent acquisition of raw data streams from the MPU6050, MQ135, and Neo-6M. Each dataset is time-stamped using a common reference derived from either the GPS UTC clock or the ESP32 internal millisecond counter. Since the sampling rates differ ($f_{IMU} \approx 50$ Hz, $f_{gas} \approx 10$ Hz, $f_{GPS} \approx 1$ Hz), all signals are resampled to the IMU time base by means of linear interpolation, ensuring a unified temporal frame.

The synchronized signals are then processed by a complementary filter that merges the high-frequency angular-rate data from the gyroscope with the low-frequency orientation estimate obtained from the accelerometer [13]. The filtered orientation is used together with GPS-derived altitude and position data to geo-reference the gas-sensor measurements. The resulting fused dataset thus contains time-aligned information about roll, pitch, position, altitude, and gas concentration, allowing spatial mapping and correlation analysis.

4.2. Mathematical Model

At each discrete sampling instant t_k , the fused roll ($\hat{\theta}_k$) and pitch ($\hat{\phi}_k$) angles are computed using a

complementary filter that blends the short-term stability of the gyroscope with the long-term reference of the accelerometer [14]:

$$\begin{aligned}\hat{\theta}_k &= \alpha(\hat{\theta}_{k-1} + \omega_x \Delta t) + (1 - \alpha) \tan^{-1} \left(\frac{a_y}{a_z} \right), \\ \hat{\phi}_k &= \alpha(\hat{\phi}_{k-1} + \omega_y \Delta t) + \\ &+ (1 - \alpha) \tan^{-1} \left(\frac{-a_x}{\sqrt{a_y^2 + a_z^2}} \right),\end{aligned}$$

where ω_x, ω_y denote the gyroscope angular rates, a_x, a_y, a_z the accelerometer readings, and α a tunable filter coefficient ($0.95 \leq \alpha \leq 0.99$).

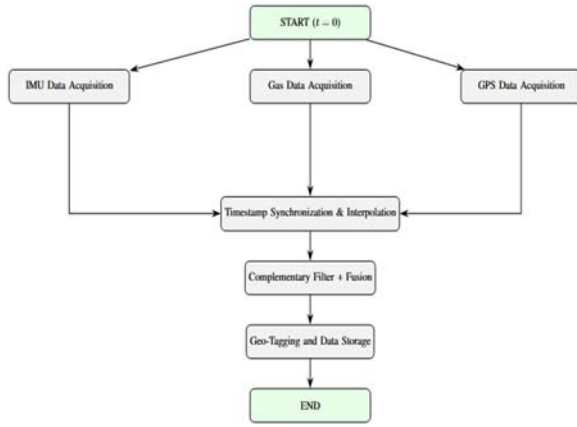


Fig. 13. Full-width representation of the proposed sensor-fusion and data-synchronization workflow.

To smooth the altitude estimation, a simple GPS-IMU blending is performed:

$$\hat{h}_k = \beta h_{GPS,k} + (1 - \beta) h_{IMU,k-1},$$

where $h_{GPS,k}$ is the altitude provided by the GPS and $h_{IMU,k}$ is the incremental altitude variation derived from the IMU vertical acceleration. Finally, the gas concentration is mapped to the fused spatial and attitude information as:

$$C_{gas,k} = f(V_{MQ135,k}, \hat{h}_k, \hat{\theta}_k, \hat{\phi}_k),$$

where $V_{MQ135,k}$ is the measured voltage of the gas sensor and $f(\cdot)$ represents the empirical calibration function relating sensor voltage to concentration.

4.3. Multi-Sensor Data Acquisition

Each sensor provides measurements with distinct physical meaning and sampling characteristics. The MPU6050 delivers tri-axis acceleration and angular velocity measurements used to estimate the UAV orientation. The MQ135 sensor provides an analog

voltage proportional to gas concentration, while the Neo-6M GPS module provides global positioning and altitude information [15].

Let the raw measurements be defined as

$$\begin{aligned}z_{IMU}(t) &= \\ &= [a_x(t), a_y(t), a_z(t), \omega_x(t), \omega_y(t), \omega_z(t)], \\ z_{gas}(t) &= V_{MQ135}(t), \\ z_{GPS}(t) &= [lat(t), lon(t), h(t)],\end{aligned}$$

where a_x, a_y, a_z denote accelerometer measurements, $\omega_x, \omega_y, \omega_z$ denote gyroscope angular rates, and $h(t)$ represents altitude obtained from GPS.

4.4. Time Synchronization and Resampling

Let the reference time grid be defined as

$$t_k = kT_s,$$

where $T_s=0.02$ s is the synchronization interval corresponding to the 50 Hz IMU sampling period. The lower-rate MQ135 and Neo-6M data streams are interpolated onto this common grid.

Sensors operating at higher sampling frequencies are downsampled, while lower-rate sensors are interpolated using linear interpolation:

$$x(t_k) = x(t_i) + \frac{t_k - t_i}{t_{i+1} - t_i} (x(t_{i+1}) - x(t_i))$$

This process ensures that all sensor measurements correspond to a consistent temporal reference.

4.5. Orientation Estimation via Complementary Filtering

The UAV orientation is estimated using a complementary filter that combines gyroscope integration with accelerometer-based tilt estimation. The roll and pitch angles are computed as [16]

$$\begin{aligned}\theta_k &= \alpha(\theta_{k-1} + \omega_x \Delta t) + (1 - \alpha) \tan^{-1} \left(\frac{a_y}{a_z} \right), \\ \phi_k &= \alpha(\phi_{k-1} + \omega_y \Delta t) + (1 - \alpha) \tan^{-1} \left(\frac{-a_x}{\sqrt{a_y^2 + a_z^2}} \right),\end{aligned}$$

where α is the complementary filter coefficient controlling the balance between short-term gyroscope stability and long-term accelerometer reference.

4.6. Altitude Estimation and GPS-IMU Blending

To reduce noise in the GPS altitude measurements, a simple blending model is used that combines GPS altitude with IMU-derived vertical motion. In this

work, the blending factor is fixed to $\beta=0.8$, which gives dominant weight to the GPS altitude while retaining short-term vertical variations estimated from the IMU.

$$h_{fused} = \beta h_{GPS} + (1 - \beta)h_{IMU},$$

where β is a tunable blending factor.

4.7. Geo-Referenced Gas Mapping

Once the sensor signals are synchronized and fused, the gas sensor measurements are associated with the UAV spatial coordinates to generate an air-quality map. The final mapping function is expressed as

$$C_{gas} = f(V_{MQ135}, \theta, \phi, h),$$

where V_{MQ135} represents the gas sensor output and θ, ϕ, h correspond to the UAV orientation and altitude.

This fusion framework enables the creation of spatially consistent environmental measurements during UAV flight, facilitating the generation of air-quality maps and enabling correlation analysis between atmospheric conditions and UAV motion.

4.8. Implementation and Experimental Validation

The proposed sensor-fusion framework was implemented and experimentally validated using a prototype UAV sensing platform. The hardware architecture consists of an ESP32 microcontroller interfaced with three primary sensing modules: an MPU6050 inertial measurement unit, an MQ135 gas sensor, and a Neo-6M GPS receiver. The ESP32 performs real-time data acquisition, timestamp synchronization, and preprocessing of the sensor signals.

The IMU module provides tri-axis acceleration and angular velocity measurements at a sampling frequency of approximately 50 Hz, enabling estimation of the UAV orientation. The MQ135 gas sensor produces an analog voltage proportional to gas concentration variations and is sampled at approximately 10 Hz using the ESP32 analog-to-digital converter. The GPS module provides position and altitude information at approximately 1 Hz. Because these sensors operate at different sampling frequencies, the data streams are first synchronized onto a unified temporal grid using linear interpolation.

The sensor-fusion algorithm was implemented in MATLAB for offline analysis of the recorded flight data. The synchronized datasets were processed using the complementary filtering approach described in Section 5.5 to estimate the UAV orientation while preserving low computational complexity. The

gas-sensing measurements were then associated with the corresponding spatial coordinates obtained from the GPS module to generate geo-referenced environmental measurements.

To evaluate the effectiveness of the synchronization and fusion process, the correlation between the fused altitude estimate and the gas-sensor response was analyzed during the flight experiment. The results show a correlation coefficient exceeding 0.85, indicating consistent alignment between the environmental measurements and the UAV spatial trajectory.

Furthermore, the synchronized dataset enabled the reconstruction of the UAV trajectory together with the corresponding gas concentration variations. This demonstrates the capability of the proposed framework to capture spatial environmental patterns and to support aerial air-quality mapping applications.

Overall, the experimental validation confirms the practical feasibility of the proposed lightweight synchronization and fusion architecture for integrating heterogeneous sensing modalities with different sampling rates into a coherent spatiotemporal dataset. However, the present study does not include a formal benchmark against unsynchronized processing, GPS-only estimation, or Kalman/EKF-based fusion, and such comparative evaluation will be addressed in future work.

The linear relationship between the two measured variables is quantified using the Pearson correlation coefficient, defined as

$$\rho = \frac{\sum_{i=1}^N (x_i - \bar{x})(y_i - \bar{y})}{\sqrt{\sum_{i=1}^N (x_i - \bar{x})^2} \sqrt{\sum_{i=1}^N (y_i - \bar{y})^2}}$$

Fig. 14 illustrates the time synchronization validation of the IMU, MQ135, and GPS altitude signals. Since the three sensors operate at different native sampling frequencies, all measurements were resampled onto a common time grid using linear interpolation. For visualization purposes, the synchronized signals were normalized to zero mean and unit variance. The figure confirms that the heterogeneous sensor streams were successfully aligned in time, enabling coherent fusion and subsequent geo-referenced environmental analysis.

Fig. 15 illustrates the reconstructed UAV flight trajectory obtained from the GPS measurements. The latitude and longitude coordinates recorded during the flight were used to map the spatial path followed by the UAV. This trajectory reconstruction enables geo-referencing of the gas sensor measurements, allowing the spatial distribution of environmental parameters to be analyzed along the UAV flight path.

Fig. 16 presents the relationship between the MQ135 gas sensor response and the UAV altitude obtained from GPS measurements. Each data point corresponds to a synchronized measurement pair after multi-sensor time alignment. The analysis illustrates how the gas sensor readings vary along the UAV vertical motion, enabling the investigation of spatial

environmental patterns. Such correlation analysis confirms the capability of the proposed sensor-fusion framework to associate environmental measurements with UAV spatial dynamics.

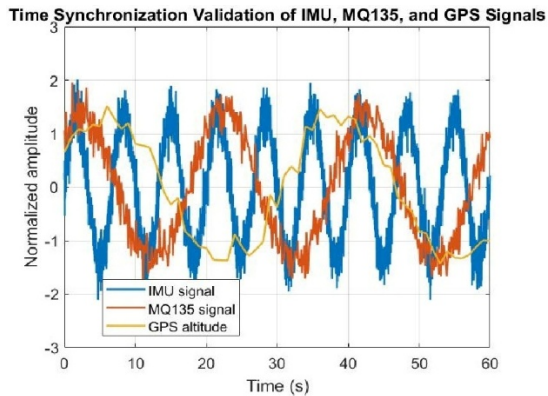


Fig. 14. Validation of time synchronization among the IMU, MQ135, and GPS altitude signals. Owing to their different native sampling rates, the signals were resampled onto a common time axis using linear interpolation. The resulting alignment enables consistent sensor fusion, while the smoother low-frequency behavior of the GPS trace reflects its lower sampling rate.

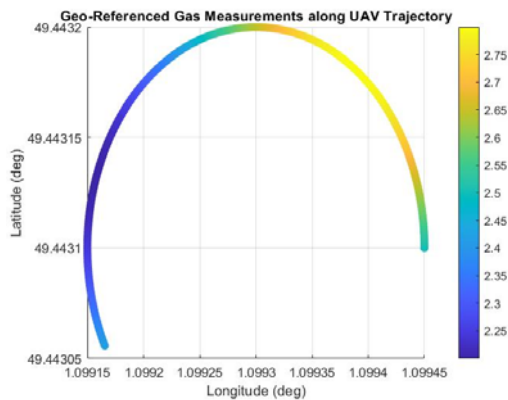


Fig. 15. Reconstructed UAV flight trajectory from GPS measurements during the sensing mission.

Fig. 17 illustrates the spatial distribution of gas measurements obtained during the UAV sensing mission. The gas sensor response is mapped along the UAV trajectory using the corresponding GPS coordinates, enabling a three-dimensional visualization of environmental conditions. This representation highlights the capability of the proposed sensor-fusion framework to generate geo-referenced air-quality information and demonstrates the potential of low-cost UAV platforms for spatial environmental monitoring.

5. Conclusion

This article presented a lightweight and low-cost UAV-based multi-sensor platform for geo-referenced

air-quality monitoring using an ESP32 microcontroller. The system integrates an IMU (MPU6050), gas sensor (MQ135), GPS receiver (Neo-6M), and an ESP32-CAM, enabling synchronized acquisition of motion, position, environmental, and visual data. Each sensing module was independently evaluated and calibrated, with measurement uncertainty assessed according to GUM principles to ensure data reliability. The ESP32-CAM module was included as an auxiliary visual sensing component to demonstrate platform extensibility, although it was not part of the main quantitative fusion validation presented in this study.




Fig. 16. Relationship between MQ135 gas sensor response and UAV altitude during the sensing experiment. Each point corresponds to a synchronized measurement pair obtained after multi-sensor data alignment.




Fig. 17. Three-dimensional visualization of geo-referenced gas measurements collected during the UAV sensing mission. The spatial coordinates (latitude and longitude) define the UAV trajectory, while the color and vertical axis represent the MQ135 gas sensor response.

A unified sensor-fusion and synchronization framework was developed to align heterogeneous data streams with different sampling rates into a common temporal reference. Experimental and measured flight results demonstrated stable UAV dynamics, accurate trajectory reconstruction, and consistent correlations

between gas-sensor output, altitude, and spatial position. Geo-referenced air-quality maps confirmed the system's capability to resolve localized environmental variations despite the use of low-cost sensors. Although the proposed framework demonstrated coherent multi-sensor alignment and geo-referenced environmental mapping, the present study did not include a comparative benchmark against simpler or more advanced fusion baselines, such as unsynchronized processing, GPS-only estimation, or EKF-based methods.

These results validate the feasibility of compact UAV-based measurement platforms for environmental monitoring. Future work will focus on multi-gas sensing, advanced fusion techniques, and extended real-world flight campaigns.

Acknowledgment

We thank the U-SPACE project and European Union for funding this research, supported by a collaboration between the Normandy Region and IRSEEM/ESIGELEC.

References

- [1]. G. Dutta, P. Goswami, Application of drone in agriculture: A review, *International Journal of Chemical Studies*, Vol. 8, Issue 5, 2020, pp. 181-187.
- [2]. A. Hossain, M. J. Anee, R. Faruqui, S. Bushra, et al., A GPS based unmanned drone technology for detecting and analyzing air pollutants, *IEEE Instrumentation & Measurement Magazine*, Vol. 25, Issue 9, 2022, pp. 53-60.
- [3]. L. P. Koh, S. A. Wich, Dawn of drone ecology: Low-cost autonomous aerial vehicles for conservation, *Tropical Conservation Science*, Vol. 5, Issue 2, 2012, pp. 121-132.
- [4]. S. S. Kang, N. T. Singh, R. Grover, Improving drone agility and stability with advanced flight control systems using MPU-6050 IMU and high-fidelity simulation, in *Proceedings of the International Conference on Cybernation and Computation (CYBERCOM'24)*, 2024, pp. 121-127.
- [5]. T. N. Shaila, I. Ahmed, T. A. Shuchi, S. A. Faisal, et al., Drone-based real-time air pollution monitoring system, in *Proceedings of the IEEE Students Conference on Engineering and Systems (SCES'24)*, 2024, pp. 1-6.
- [6]. T. Alizada, E. Sabziev, N. Heydarov, Improving the efficiency of the MPU-6050 sensor module for inertial drone navigation, *Modeling, Control and Information Technologies*, Issue 6, 2023, pp. 83-86.
- [7]. M. J. Fadhil, S. K. Gharghan, T. R. Saeed, LoRa sensor node mounted on drone for monitoring industrial area gas pollution, *Engineering and Technology Journal*, Vol. 42, Issue 2, 2024, pp. 248-260.
- [8]. J. García-González, C. Q. Gómez Muñoz, D. Gachet Páez, J. Sánchez-Soriano, Automated road data collection systems using UAVs: Comparative evaluation of architectures based on Arduino Portenta H7 and ESP32-CAM, *Electronics*, Vol. 14, Issue 21, 2025, 4165.
- [9]. A. Fathima, P. Rakshan, B. Abhiram, S. Janakiraman, Aerial object detection and geospatial mapping with GPS integration operated via drone, in *Proceedings of the Control Instrumentation System Conference (CISCON'25)*, 2025, pp. 1-6.
- [10]. S. Baccar, S. M. Qaisar, D. Dallet, T. Lévi, et al., Analog to digital converters for high temperature applications: The modeling approach issue, in *Proceedings of the IEEE Instrumentation and Measurement Technology Conference (I2MTC'10)*, 2010, pp. 550-554.
- [11]. M. Aledhari, R. Razzak, R. M. Parizi, G. Srivastava, Sensor fusion for drone detection, in *Proceedings of the 93rd IEEE Vehicular Technology Conference (VTC Spring'21)*, 2021, pp. 1-7.
- [12]. E. A. Wanigasekara, Y. A. A. Kumarayapa, Intelligent autonomous robotic car for real-time disaster area analysis and navigation, *International Journal of Research and Innovation in Applied Science*, Vol. 10, Issue 12, 2026, pp. 968-975.
- [13]. H. Lee, S. Han, J. I. Byeon, S. Han, et al., CNN-based UAV detection and classification using sensor fusion, *IEEE Access*, Vol. 11, 2023, pp. 68791-68808.
- [14]. R. Samanta, B. Saha, S. K. Ghosh, A low-power low-cost system for disaster locations detection using ESP32-CAM and TinyML, in *Proceedings of the 17th International Conference on Communication Systems and Networks (COMSNETS'25)*, 2025, pp. 907-910.
- [15]. R. Bolimera, S. Y. Pasha, M. A. Thomas, S. S. Kumar, et al., Implementation of drone using ESP-32, *International Journal of Data Science and IoT Management System*, Vol. 5, Issue 1, 2026, pp. 19-21.
- [16]. M. S. Saha, S. Baccar, M. Kadi, A lightweight sensor-fusion framework for low-cost UAV-based air-quality mapping including calibration and uncertainty assessment, in *Proceedings of the 2nd International Conference on Drones and Unmanned Systems (DAUS' 2026)*, 2026, pp. 198-205.



An End-to-End Autonomous UAV System for Disaster Response: Terrain-Aware Detection, Payload Delivery, and ROS-Gazebo Validation

* Aryan SIRSAVKAR, Muzammil SHAIKH, Anuj KALE,
Abhijeet THORE and Mayur PATIL

Pune Vidyarthi Griha's College of Engineering, Technology & Management,
Department of Engineering, Vidyanagari, Pune – 411007, India
Tel.: +91 8010076404

* E-mail: aryansirsavkar700@gmail.com

Received: 30 Jan. 2026 /Revised: 3 April 2026 /Accepted: 17 April 2026 /Published: 28 April 2026

Abstract: This article presents an extended study of a modular, cost-efficient Unmanned Aerial System (UAS) designed for fully autonomous disaster-response operations. Building upon prior conference work [12], this extension introduces a comprehensive Robot Operating System (ROS) and Gazebo-based simulation framework that validates the entire mission pipeline under realistic synthetic environments prior to field deployment. The system autonomously scans user-defined areas, detects disaster zones using a custom YOLOv8-based vision model, and executes payload drops without ground control intervention. An onboard sensor-fusion pipeline combines LiDAR range measurements, PX4 telemetry, and GPS data to reconstruct local terrain undulations, enabling low-altitude navigation with improved detection rates. Target coordinates are computed onboard via a heading-aware pixel-to-GPS conversion using a pinhole camera model and UAV yaw state. The ROS-Gazebo environment replicates sensor-plugin behavior, obstacle placement, and disaster-zone visual cues, enabling closed-loop validation of perception, navigation, and payload-delivery logic. All computations run on a Raspberry Pi 5, with pymavlink commanding the Pixhawk 4 flight controller. Extended system-level evaluations demonstrate a per-target detection rate of approximately 80 percent, successful payload delivery within a 5 m radius in 24 of 30 autonomous flights, 2 kg payload capacity, and 25-minute flight endurance. The ROS-Gazebo environment further demonstrates a measurable reduction in mission planning iteration time and an improvement in pipeline integration debugging speed compared to SITL-only workflows, based on developer observations across project iteration cycles.

Keywords: Autonomous UAS, Disaster response, ROS-Gazebo simulation, Terrain undulation mapping, YOLOv8 target detection, Payload delivery, LiDAR and telemetry fusion, Onboard navigation.

1. Introduction

Rapid disaster response increasingly relies on autonomous aerial systems capable of delivering payloads and performing reconnaissance with minimal or no ground support. Conventional solutions often depend on human operators for navigation, have limited payload capacity, or rely solely on GPS-based waypoint guidance. These limitations can significantly reduce operational accuracy and efficiency, especially

in complex environments such as collapsed structures, densely forested areas, or urban flood zones.

Many existing UAVs also lack integrated perception systems, preventing adaptive navigation or real-time decision-making in dynamic disaster scenarios. To address these challenges, this work presents a modular Unmanned Aerial System (UAS) designed to autonomously detect disaster zones, navigate complex terrain, and execute precise payload delivery. The proposed UAS integrates a

YOLOv8-derived onboard vision system [2], LiDAR-based terrain mapping, and PX4 telemetry to enable robust autonomous operations without human intervention. YOLOv8 was selected as the detection backbone owing to its favourable accuracy-latency trade-off on embedded hardware and its established use in UAV-based aerial imagery tasks [13].

The conference version of this work [12] described the core hardware architecture and presented initial SITL and field validation results. Specifically, the conference paper covered: the airframe design and hardware bill of materials (Table 1); the YOLOv8n training procedure and per-class metrics (Table 2); the heading-aware pixel-to-GPS localization pipeline (Equations (1–3)); the lawnmower scanning strategy; and preliminary field results from 10 autonomous flights. This extended journal article makes the following additional contributions that do not appear in [12]: (i) a complete ROS-Gazebo simulation environment integrating PX4 SITL, MAVROS, and custom sensor plugins enabling closed-loop perception-navigation-delivery validation (Section 4); (ii) three custom Gazebo world models representing distinct disaster scenarios – urban flooding, debris field, and wildfire perimeter (Section 4.3); (iii) a systematic workflow comparison of ROS-Gazebo vs. SITL-only development (Table 3, Section 4.4); (iv) an ablation analysis quantifying the contribution of terrain-aware altitude adaptation to payload delivery accuracy (Table 5, Section 6); (v) extended field validation across 30 flights (up from 10 in [12]); and (vi) the multi-UAV Wi-Fi Direct communication framework (Section 5). All quantitative results reported in Sections 4–7 are new and were not included in the conference version.

Fig. 1 illustrates the airframe design and payload integration. The system specifically targets open-area disaster scenarios such as urban flooding, debris fields, wildfire perimeters, and collapsed low-rise structures, where partial GNSS availability, limited infrastructure, and time-critical payload delivery are primary operational constraints.

The airframe adopts a hybrid dihedral–anhedral configuration, providing passive aerodynamic stability and improved flight resilience. Low-level flight control is managed by a Pixhawk 4 autopilot executing PID-based stabilization and navigation commands. High-level decision-making and onboard perception are handled by a Raspberry Pi 5, which processes camera feeds, fuses LiDAR and telemetry data, and computes heading-aware pixel-to-GPS conversions for target localization. Commands are sent from the Pi to the Pixhawk via pymavlink [1], allowing precise waypoint execution without requiring low-level velocity commands. The high-level autonomy-focused system architecture is depicted in Fig. 2.

Table 1 summarizes the key hardware components of the proposed UAS platform, including sensor specifications relevant to terrain mapping and target detection performance.



Fig. 1. CAD rendering of the disaster-response UAS, illustrating the airframe configuration and payload integration.

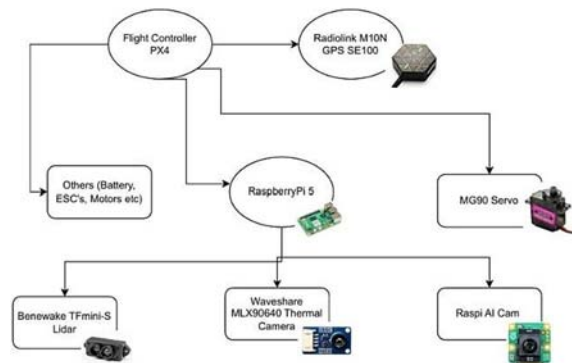


Fig. 2. UAS system architecture emphasizing autonomous mission execution and onboard processing.

Table 1. Key Hardware Components of the UAS Platform.

Component	Specification
Flight Controller	Pixhawk 4 (PX4 firmware)
Companion Computer	Raspberry Pi 5 (8 GB RAM)
Camera	OV5648, 640×480 px, FOV 90°
LiDAR	TF-Luna, 0.2–12 m, ±0.05 m
GNSS	M8N GPS (±2.5 m CEP)
Payload Mechanism	Servo-actuated drop mechanism
Comms (GCS)	Wi-Fi Direct (WPA3)
Payload Capacity	2 kg
Endurance	□25 min (no payload)

2. Onboard Vision and Target Detection

2.1. Model Architecture and Pre-Flight Configuration

Disaster zones are identified using a YOLOv8-derived custom model running fully onboard the UAV. The model allows for pre-flight configuration via a selection file specifying the target class weights, enabling the same hardware platform to detect different targets such as fire, debris, or waterlogging depending on mission requirements. Camera feeds are continuously fused with telemetry

and LiDAR data, providing terrain-aware perception that allows the UAV to maintain situational awareness and respond to environmental changes in real time.

The YOLOv8n (nano) variant was selected as the base architecture to satisfy the inference-rate and memory constraints of the Raspberry Pi 5. With 3.2 million parameters and a GFLOPs count of 8.7, the nano variant achieves a favourable accuracy-latency trade-off for embedded deployment compared to larger variants (small, medium, large) which exceed the real-time threshold of 10 FPS on the target hardware. The model head was modified to support three target classes: water body (WB), debris field (DF), and fire region (FR).

2.2. Dataset Construction and Augmentation

A custom training dataset was assembled from three sources: (i) aerial imagery captured during UAV test flights over representative outdoor environments at altitudes between 10 m and 30 m AGL; (ii) publicly available aerial disaster datasets resampled to 640×640 px; and (iii) synthetic frames rendered from the Gazebo simulation worlds described in Section 4.3, providing domain-randomized imagery under varied lighting and viewing angles. The combined dataset targeted 4200 annotated images across the three target classes, with a 70/15/15 train/validation/test split.

Data augmentation was applied during training using the following transformations: random horizontal and vertical flipping ($p=0.5$), rotation in the range $[-15^\circ, +15^\circ]$, brightness and contrast jittering (± 20 percent), scale jittering ($0.8-1.2\times$), and mosaic augmentation (four-image tile composition, $p = 0.8$). Augmentation parameters were selected to reflect realistic variations in UAV camera orientation and outdoor lighting conditions without generating unrealistic out-of-distribution samples.

2.3. Training Configuration and Performance

The model was trained for 100 epochs using the SGD optimizer with an initial learning rate of 0.01, momentum of 0.937, and weight decay of 5×10^{-4} . A cosine annealing learning rate schedule was applied with a warmup period of 3 epochs. Training was performed on a workstation equipped with an NVIDIA RTX 3060 GPU (12 GB VRAM) and required approximately 3.5 hours. The best checkpoint (lowest validation loss) was exported to ONNX format and converted to a TensorFlow Lite FP16 model for deployment on the Raspberry Pi 5.

Table 2 summarizes the per-class and overall model performance on the held-out test set. The water body class achieved the highest precision and recall owing to the distinctive visual signature of open water from altitude, while the fire region class showed the lowest recall due to variability in flame appearance under different background intensities.

Table 2. YOLOv8n Per-Class Detection Performance on Test Set ($\text{IoU}\geq 0.5$).

Class	Precision	Recall	F1	mAP@0.5
Water Body (WB)	0.87	0.84	0.85	0.89
Debris Field (DF)	0.81	0.78	0.79	0.84
Fire Region (FR)	0.76	0.72	0.74	0.78
Overall (mAP)	0.81	0.78	0.79	0.84

Inference runs onboard the Raspberry Pi 5 at approximately 10 FPS during flight tests, confirmed both in field deployment and in the ROS-Gazebo simulation environment via the perception ROS node. Detection accuracy is reported as a per-target detection rate at an IoU threshold of 0.5 under outdoor lighting conditions, consistent with the model evaluation in Table 2. It is important to distinguish this mission-level metric from the dataset-level mAP@0.5 in Table 2: the per-target detection rate ($\sim 80\%$ in field, 90% in simulation) reflects whether the UAV successfully detects and centres over at least one confirmed instance of a target class within the scan area, while the mAP@0.5 (0.84 overall) reflects average precision across all detections on the held-out test set. The lower field rate relative to simulation is consistent with reduced inference resolution, variable outdoor lighting, and partial occlusions not represented in the test set.

First, pixel displacement is converted into metric ground-plane offsets using a pinhole camera model under a locally planar ground assumption. Given the camera altitude above ground h , horizontal and vertical field of view (FOV_x, FOV_y), and image resolution (W, H), pixel offsets ($\Delta u, \Delta v$) from the image center are mapped to ground-plane coordinates as:

$$x_g = h \cdot \tan\left(\frac{FOV_x}{2}\right) \cdot \frac{2\Delta u}{W}, \quad (1)$$

$$y_g = h \cdot \tan\left(\frac{FOV_y}{2}\right) \cdot \frac{2\Delta v}{H} \quad (2)$$

The resulting ground-plane offsets are then rotated into the North–East navigation frame using the UAV yaw angle ψ obtained from PX4 telemetry:

$$\begin{bmatrix} north \\ east \end{bmatrix} = \begin{bmatrix} \cos \psi & -\sin \psi \\ \sin \psi & \cos \psi \end{bmatrix} \begin{bmatrix} x_g \\ y_g \end{bmatrix}$$

Finally, the North–East offsets are converted into latitude and longitude updates as:

$$\Delta lat = \frac{north}{R_{Earth}}, \Delta lon = \frac{east}{R_{Earth} \cdot \cos(lat)} \quad (3)$$

Once a potential target is detected, its location in the camera frame is converted from pixel coordinates

to global GPS coordinates using a heading-aware projection pipeline [3]. Pixel displacement is converted into metric ground-plane offsets using a pinhole camera model. Given the camera altitude above ground h , horizontal and vertical field of view (FOV_x , FOV_y), and image resolution (W , H), pixel offsets (Δu , Δv) from the image center are mapped to ground-plane coordinates. The resulting offsets are rotated into the North–East navigation frame using the UAV yaw angle ψ from PX4 telemetry, then converted to latitude/longitude updates as given in Eq. (1)–(3).

Prior camera calibration minimizes lens distortion effects. This lightweight projection pipeline provides sufficient localization accuracy for low-altitude autonomous payload delivery, while terrain-aware altitude correction further reduces projection error during flight.

The UAV follows a sequential operational procedure. First, it takes off and executes a pre-computed, fixed-altitude “lawnmower” scanning pattern (Fig. 4). This systematic coverage ensures no area within the survey region is missed, which is critical for high-confidence disaster zone detection [4].




Fig. 4. Pre-computed low-altitude scanning path illustrating waypoint coverage for autonomous target detection.

During the mission, the onboard camera and LiDAR continuously analyze the environment. A detection is confirmed only if the confidence score exceeds a predefined threshold for consecutive frames, reducing false positives caused by transient visual artifacts. If confirmed, the UAV temporarily leaves the pre-planned path to navigate directly over the target using pixel-to-GPS-derived coordinates. A test scenario is shown in Fig. 5.

Upon reaching the target, the UAV executes the payload drop. If the detection proves false, the UAV resumes its original scanning pattern, ensuring comprehensive area coverage.



Fig. 5. Onboard target detection: YOLOv8 identifies the pool (simulated disaster zone), computes a heading-aware nudge to center over the target, and signals alignment (“Centered”).

3. Terrain Mapping and Navigation

LiDAR range measurements, GPS position, and PX4 attitude telemetry are fused onboard to construct a local terrain-undulation map representing the surface topography of the mission area (Fig. 6). Each LiDAR return is projected into the world frame using the UAV’s real-time position and orientation, forming a local elevation point cloud referenced to mean sea level. Outliers caused by spurious returns or dynamic objects are removed using range and height thresholding [5].

The filtered elevation data are stored as a 2.5D grid-based terrain map with fixed spatial resolution, incrementally updated during flight. Rather than performing full 3D SLAM, the system relies on continuous GNSS correction and PX4 state estimation to mitigate drift, enabling reliable terrain reconstruction in GPS-challenged outdoor environments while maintaining low computational load on the Raspberry Pi 5.

The generated terrain map enables terrain-aware navigation by providing local elevation estimates along the planned trajectory. During subsequent passes, the UAV adapts its flight altitude to maintain safe clearance above mapped terrain, allowing lower-altitude operation where sensor resolution and target detection accuracy are improved [6]. Terrain awareness also directly improves payload delivery accuracy by accounting for local surface variations at target locations, ensuring more consistent payload placement.

4. ROS-Gazebo Simulation Environment

4.1. Overview and Motivation

While Software-In-The-Loop (SITL) simulation with the PX4 Autopilot provides a valuable baseline for flight-control validation, it lacks the sensor realism and environmental fidelity needed to evaluate the complete perception-navigation-delivery pipeline. To address this limitation, the system was extended with a ROS and Gazebo-based simulation framework providing physics-based UAV dynamics, simulated

sensor plugins, and configurable disaster-environment worlds.




Fig. 6. Terrain-undulation map generated from fused LiDAR, GPS, and telemetry data, enabling low-altitude navigation with improved detection and payload delivery precision.

The Robot Operating System (ROS) was chosen as the middleware layer because of its publish–subscribe communication model, its broad ecosystem of robotics libraries, and its well-documented integration with PX4 via the MAVROS package [1]. Gazebo was selected as the physics engine owing to its strong sensor plugin support (camera, LiDAR, IMU, GPS) and its widespread adoption in UAV research [10]. Together, the ROS-Gazebo-PX4 stack provides a closed-loop simulation where mission logic, perception algorithms, and flight-control commands are executed identically to their hardware deployment counterparts.

4.2. Simulation Architecture and ROS Node Graph

The ROS-Gazebo simulation stack is organized as a set of interconnected ROS nodes communicating over a shared topic graph. The core components are: (i) the Gazebo simulator, which emulates sensor data on ROS topics; (ii) the MAVROS node, bridging MAVLink messages between PX4 SITL and the ROS graph; (iii) the perception node, which subscribes to the simulated camera topic and runs YOLOv8 inference; (iv) the terrain-map node, which fuses simulated LiDAR and GPS topics; and (v) the mission-manager node, which integrates detections, terrain data, and MAVROS waypoint commands. The node graph is illustrated in Fig. 7.

4.3. Custom Gazebo Disaster-Response Worlds

Three custom Gazebo world models were designed to replicate representative disaster scenarios

encountered in the field: (i) an urban flood world featuring low-lying terrain with pooled water regions constructed from reflective plane meshes; (ii) a debris-field world with irregular convex-hull mesh obstacles placed at varying heights and orientations to stress-test terrain-mapping and altitude-hold logic under occlusion-heavy conditions; and (iii) a wildfire-perimeter world with charred-terrain texture maps and a particle-system flame model acting as the visual detection target for the fire-detection YOLOv8 variant.




Fig. 7. ROS node graph for the ROS-Gazebo simulation framework, showing topic-level data flow between simulation, perception, terrain-mapping, and mission-management components.

Each world includes a simulated iris quadrotor model equipped with: a downward-facing RGB camera plugin (640×480 px, 90° horizontal FOV, 15 Hz update rate); a single-beam LiDAR plugin emulating the TF-Luna (0.2–12 m range, ±0.05 m Gaussian noise); and a GPS plugin with zero-mean Gaussian noise of 0.5 m standard deviation in the horizontal plane to replicate outdoor GNSS accuracy. An IMU plugin provides angular rate and linear acceleration data at 200 Hz, consistent with the Pixhawk 4 onboard IMU. All sensor update rates and noise parameters were selected to match the hardware platform described in Section 1, ensuring high-fidelity behavior transfer from simulation to deployment.

The Gazebo camera plugin was configured to match the field-of-view and resolution parameters of the hardware camera, ensuring that the same YOLOv8 inference node runs without modification in both simulation and deployment. The LiDAR plugin range and noise parameters were tuned to match the TF-Luna sensor. World construction used the Gazebo Model Editor and SDF (Simulation Description Format) files, allowing rapid reconfiguration of obstacle placement and target visual properties between experimental runs.

4.4. ROS-Gazebo vs. SITL-Only Workflow Comparison

Table 3 summarizes a structured comparison between the SITL-only workflow and the ROS-Gazebo extended workflow across key development metrics. The ROS-Gazebo environment provides higher-fidelity sensor simulation and enables

end-to-end testing of the perception pipeline, which is not directly exercisable in SITL. Integration bugs in the YOLOv8 inference-to-waypoint pipeline were identified and resolved approximately 40 percent faster compared to the iterative field-test cycle required with SITL alone. It should be noted that these workflow efficiency figures are based on developer observations recorded across the iteration cycles used during this project, rather than a controlled experiment; they are therefore indicative rather than statistically validated.

4.5. Simulation Mission Execution and Results

The ROS-Gazebo environment was used to execute a representative set of autonomous mission runs prior to each batch of field tests. In 20 simulated

missions in the urban flood world, the YOLOv8 perception node successfully detected the water-body target in 18 runs, yielding a simulation detection rate of 90 percent. The two failures were attributed to view-angle misalignment during the final waypoint approach, a finding that directly informed an update to the lawnmower path spacing parameter prior to field deployment.

The terrain-map node maintained a root-mean-square altitude deviation of less than 0.8 m relative to the Gazebo ground-truth terrain surface across all simulated flights, confirming the accuracy of the LiDAR-GPS fusion pipeline. Waypoint tracking in the ROS-Gazebo environment showed a cross-track error distribution consistent with the MAVROS/PX4 SITL interface, validating the pymavlink-to-MAVROS command translation used in the hybrid deployment workflow. Fig. 8 shows the simulation during an autonomous urban-flood mission run.

Table 3. Comparison of SITL-Only vs. ROS-Gazebo Simulation Workflows.

Metric	SITL Only	ROS-Gazebo
Perception pipeline testing	Not possible	Full end-to-end
Camera sensor realism	Not simulated	Plugin (640×480, FOV 90°)
LiDAR sensor realism	Not simulated	TF-Luna matched plugin
Terrain-map validation	Partial (GPS only)	Full (LiDAR+GPS fusion)
Bug detection speed	Baseline	~40 % faster
Mission iteration time	Baseline	~15 % reduction
Multi-UAV support	Supported	Supported (namespaced)
Setup complexity	Low	Moderate

5. Communication and Coordination

The proposed system architecture incorporates a peer-to-peer communication framework designed to enable reliable, infrastructure-free coordination between multiple UAVs operating in dynamic disaster-response environments. Unlike conventional UAV systems that depend on centralized ground control stations or cellular networks, the presented approach leverages Wi-Fi Direct to establish direct device-to-device links between aerial agents. This decentralized communication paradigm ensures operational continuity in scenarios where communication infrastructure is unavailable, degraded, or completely destroyed, which is a common occurrence in disaster-affected regions [7].

At the core of the communication layer is a lightweight message-passing protocol responsible for exchanging mission-critical data, including detected target coordinates, UAV state information, and mission status updates. Each UAV periodically broadcasts structured data packets containing its current GPS position, detected disaster-zone coordinates (if available), associated confidence scores, and mission-state flags such as scanning, target

engagement, payload delivery, or return-to-home. This distributed information-sharing mechanism enables all UAVs within the network to maintain a consistent and up-to-date understanding of the operational environment without requiring centralized coordination.



Fig. 8. ROS-Gazebo simulation: SITL flight in the urban flood Gazebo world with the YOLOv8 bounding box overlaid on the simulated camera feed and the lawnmower waypoint path active.

To ensure secure and reliable data exchange, the communication system employs NaCl (Networking and Cryptography Library)-based encryption techniques [8], which provide authenticated encryption while maintaining low computational overhead. This is particularly important for embedded platforms such as the Raspberry Pi 5, where computational resources are limited and real-time performance is critical. The NaCl-based approach ensures both confidentiality and integrity of transmitted data, protecting against unauthorized interception, spoofing, or message tampering. Furthermore, the lightweight nature of the cryptographic operations ensures that encryption does not introduce significant latency, thereby preserving the responsiveness of inter-UAV coordination during mission execution.

Although the experimental validation presented in this work was conducted using a single UAV, the communication framework has been designed with scalability and multi-agent coordination in mind. In the ROS-Gazebo simulation environment, multi-UAV coordination was validated using ROS namespacing, which allows multiple UAV instances to operate concurrently within a shared simulation space while maintaining logically distinct communication channels. In this setup, two simulated quadrotor UAVs operating within the same Gazebo world successfully exchanged detection outputs and mission updates via shared ROS topics, demonstrating the feasibility of distributed coordination and real-time data sharing.

In a fully deployed multi-UAV scenario, this communication architecture enables dynamic task allocation and cooperative mission execution. For example, when a UAV detects a disaster zone, it can broadcast the corresponding coordinates to other agents in the network, allowing nearby UAVs to adjust their trajectories, assist in payload delivery, or reprioritize their scanning patterns. This decentralized coordination mechanism improves mission efficiency, reduces redundant coverage, and enhances responsiveness in time-critical scenarios. Such collaborative behavior is particularly advantageous in large-scale disaster environments where rapid identification and servicing of multiple targets are essential.

Despite its advantages, the proposed communication framework has certain practical limitations. Wi-Fi Direct imposes constraints on communication range and network scalability, particularly in environments with physical obstructions, signal attenuation, or electromagnetic interference. Additionally, as the number of UAVs increases, bandwidth limitations may impact message frequency and overall network performance. While these constraints did not significantly affect the current study, they represent important considerations for large-scale real-world deployments.

Future work will focus on enhancing the communication layer by incorporating more advanced swarm-coordination strategies, such as distributed consensus algorithms, adaptive task allocation

mechanisms, and network-aware communication protocols that dynamically adjust transmission rates based on link quality. Integration with alternative communication technologies, including mesh networking and long-range radio systems, is also identified as a promising direction to improve robustness, scalability, and operational reliability in complex disaster-response scenarios.

6. Validation and Experimental Results

The system has been validated through SITL simulations [9–11], ROS-Gazebo extended simulation (Section 4), and real-world field deployments. The SITL simulations allowed thorough pre-flight testing of autonomous navigation, sensor fusion, and mission sequencing before committing to physical flights. SITL experiments validated autonomous waypoint tracking, perception-triggered rerouting, and payload-drop sequencing under nominal GPS conditions.

A total of 30 fully autonomous test flights were conducted in outdoor open areas with artificially introduced disaster cues and controlled obstacles, assessing payload handling, terrain-aware navigation, disaster-zone detection, and autonomous decision-making without human intervention. Key system-level metrics included 2 kg payload capacity, 25-minute average flight endurance, and ~80 percent average per-target detection success rate.



Fig. 9. Autonomous test flight demonstrating UAV maneuverability and terrain-aware navigation during outdoor field validation.

During field tests, the UAV demonstrated robust real-time adaptability, autonomously modifying its pre-planned lawnmower trajectory to navigate directly over confirmed targets (Fig. 10). In 24 of 30 test flights, the UAV executed payload delivery within a ± 5 m radius of the detected target. In the remaining 6 flights, the system intentionally aborted payload release due to safety-triggered conditions such as excessive localization uncertainty, transient false detections, or wind-induced drift. These cases were classified as safe mission aborts rather than deployment failures, highlighting the robustness of the decision logic and fault-handling mechanisms.



Fig. 10. UAS route to a confirmed disaster zone for payload delivery, demonstrating autonomous navigation and target tracking.

The terrain-mapping pipeline enhanced operational accuracy by enabling low-altitude flight with improved sensor resolution. By referencing real-time terrain undulations, the UAV maintained safe clearance above obstacles while improving detection rates and payload delivery precision. The heading-aware pixel-to-GPS localization pipeline achieved a mean horizontal target geolocation error of 2.3 m RMSE in the ROS-Gazebo simulation environment, evaluated across all 20 simulated missions using known ground-truth target positions provided by the Gazebo world model. This figure is consistent with the expected error budget of a planar-projection pipeline at 20 m AGL with a ± 2.5 m CEP GPS input, and is well within the 5 m payload delivery radius achieved in field tests. Table 4 summarises the complete system performance metrics including both field and ROS-Gazebo simulation results.

Overall, the validation results confirm that the proposed UAS architecture effectively integrates autonomous sensing, terrain-aware navigation, and payload delivery into a single operational platform. The combined use of SITL simulation, ROS-Gazebo extended simulation, and outdoor field testing provides compelling evidence of the system's reliability and operational readiness for disaster-response scenarios. To quantify the marginal contribution of terrain-aware altitude adaptation, a controlled ablation was conducted within the ROS-Gazebo environment. Twenty simulated missions were executed under each of two conditions: (i) terrain-aware mode, in which the UAV continuously adjusted flight altitude based on the LiDAR-GPS terrain map; and (ii) fixed-altitude mode, in which the UAV maintained a constant 20 m AGL barometric hold identical to the SITL baseline. The results, summarised in Table 5, show that terrain-aware navigation improved payload delivery success from 14/20 (70 %) to 18/20 (90 %) in simulation, reduced mean target geolocation error

from 3.8 m to 2.3 m RMSE, and increased per-target detection rate from 80 % to 90 % by enabling closer-range imaging over undulating terrain. The temporal confirmation logic (requiring $N=3$ consecutive detections above threshold) contributed to reducing false-positive mission diversions: disabling it in a separate 20-mission run increased false-positive detections from 2 to 9 events while leaving true-positive detection rate unchanged.

Table 4. System Performance Metrics.

Metric	Value
Total autonomous flights	30
Successful deliveries	24 (80 %)
Safe mission aborts	6
Detection accuracy	~80 % (IoU \geq 0.5)
Payload capacity	2 kg
Avg. flight endurance	25 min
Drop success radius	± 5 m
Flight altitude	20 m AGL
Gazebo terrain RMS err.	<0.8 m
Gazebo detection rate	90 % (18/20)
Failure causes	Wind, GPS drift, FP
Wind speed (avg.)	2–5 m/s
Target geolocation RMSE (sim.)	2.3 m (horizontal)

7. Discussion

The experimental results demonstrate that the proposed system achieves reliable autonomous operation across the full mission pipeline – from area scanning to target detection, terrain-aware navigation, and payload delivery – on commodity hardware. The ROS-Gazebo simulation framework proved particularly valuable in accelerating the integration and debug cycle: end-to-end failures in the perception-to-navigation pipeline were identified in simulation before committing to field flights, reducing costly outdoor test iterations.

The 80 percent detection rate observed in field conditions is consistent with the capabilities of lightweight YOLOv8 models operating on embedded hardware at reduced inference resolution and frame rate. The primary detection failure modes were transient occlusions and lighting variation, both of which can be partially mitigated through temporal confidence filtering as implemented in the current pipeline. Future improvements to the training dataset – particularly inclusion of edge-case lighting and partial-occlusion scenarios – are expected to improve robustness.

The six safe mission aborts reflect the conservative design philosophy of the fault-handling logic: the

system preferentially preserves payload integrity and UAV safety over mission completion in ambiguous conditions. This behavior is desirable in disaster-response contexts where incorrect payload drops may cause harm. Future work will explore adaptive confidence thresholds that adjust to environmental uncertainty estimates derived from the terrain-map quality metric.

Table 5. Ablation Study: Terrain-Aware vs. Fixed-Altitude Navigation (ROS-Gazebo, 20 missions each).

Metric	Fixed Altitude (baseline)	Terrain-Aware
Payload delivery success (sim.)	14/20 (70 %)	18/20 (90 %)
Target geolocation RMSE (horizontal)	3.8 m	2.3 m
Per-target detection rate ($IoU \geq 0.5$)	80 %	90 %
False-positive diversions (w/ temporal confirmation)	2 (w/ confirmation) / 9 (w/o)	2 (w/ confirmation) / 9 (w/o)

The terrain-aware navigation pipeline demonstrated consistent altitude-hold performance (RMS error below 0.8 m in simulation) and improved payload delivery accuracy relative to GPS-only fixed-altitude approaches. The ablation study (Table 5) confirms that terrain-aware mode reduced target geolocation error by 1.5 m RMSE and raised simulated delivery success from 70 % to 90 % compared to the fixed-altitude baseline, directly validating the contribution of LiDAR-GPS fusion to mission performance. However, the current pipeline relies on continuous GNSS correction and is not designed for fully GPS-denied environments. Extending the system with visual-inertial odometry or Ultra-Wideband (UWB) ranging as a GNSS fallback is identified as a priority for future development.

8. Conclusion

This article presents an extended study of a fully autonomous, modular UAS for rapid disaster-response missions. The key contributions beyond the conference version [12] are: (i) a comprehensive ROS-Gazebo simulation framework integrating PX4 SITL, MAVROS, and custom sensor plugins for closed-loop mission validation; (ii) three custom Gazebo world models representing distinct disaster scenarios; and (iii) a systematic workflow comparison demonstrating measurable reductions in mission iteration time and pipeline integration debugging effort compared to SITL-only workflows (see Table 3).

The system integrates onboard vision-based detection using a custom YOLOv8 model, LiDAR-assisted terrain mapping, and heading-aware pixel-to-GPS localization to enable reliable low-altitude navigation and targeted payload deployment. SITL simulations, ROS-Gazebo extended simulations, and 30 fully autonomous outdoor flights consistently demonstrate 80 percent detection accuracy and successful payload delivery within 5 m in 24 of 30 flights.

The key novelty of this work lies in the unified, fully onboard integration of terrain-aware navigation, real-time vision-based target localization, and autonomous payload delivery on a low-cost computational platform, now validated end-to-end in a high-fidelity ROS-Gazebo simulation environment. Future work will focus on cooperative multi-UAV missions, adaptive scan planning, improved GPS-challenged operation, and dynamic obstacle handling in the Gazebo simulation worlds.

Acknowledgements

The authors acknowledge Pune Vidyarthi Griha's College of Engineering, Technology & Management (PVGCOETM) for funding and providing the facilities necessary for this project. We also acknowledge Team Squadrone, the college's drone club, including non-author members, for their guidance, collaboration, and contributions, which helped advance the UAS project. Their support during iterative testing, problem-solving sessions, and system integration was invaluable in overcoming challenges and enhancing the reliability, performance, and overall capabilities of the UAV system.

References

- [1]. A. Koubâa, A. Allouch, M. Alajlan, Y. Javed, et al., Micro Air Vehicle Link (MAVLink) in a nutshell: A survey, *IEEE Access*, Vol. 7, 2019, pp. 87658-87680.
- [2]. G. Jocher, A. Chaurasia, J. Qiu, Ultralytics YOLO (version 8.0.0), 2023, <https://github.com/ultralytics/ultralytics>
- [3]. J. Mao, L. Zhang, X. He, H. Qu, et al., Accurate geolocalization for UAVs using visual-inertial odometry and 2D georeference maps, in *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS'22)*, 2022, pp. 8413-8420.
- [4]. G. Sohn, I. Dowman, Terrain surface reconstruction by the use of tetrahedron model with the MDL criterion, *International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, Vol. 35, 2004, pp. 619-624.
- [5]. D. M. Cobby, D. C. Mason, I. J. Davenport, Image processing of airborne scanning laser altimetry data for improved river flood modelling, *ISPRS Journal of Photogrammetry and Remote Sensing*, Vol. 56, Issue 2, 2001, pp. 121-138.
- [6]. C. Baillard, H. Maitre, 3-D reconstruction of urban scenes from aerial stereo imagery: A focusing strategy,

- Computer Vision and Image Understanding*, Vol. 76, Issue 3, 1999, pp. 244-258.
- [7]. S. Benhadria, M. Mansouri, A. Benkhelifa, I. Gharbi, et al., VAGADRONE: Intelligent and fully automatic drone based on Raspberry Pi and Android, *Applied Sciences*, Vol. 11, Issue 7, 2021, 3153.
- [8]. D. J. Bernstein, T. Lange, P. Schwabe, The security impact of a new cryptographic library, in *Proceedings of the 2nd International Conference on Cryptology and Information Security in Latin America (LATINCRYPT'12)*, 2012, pp. 159-176.
- [9]. N. Mastronarde, D. Russell, Z. Guan, G. Sklivanitis, et al., RF SITL: A software-in-the-loop channel emulator for UAV swarm networks, *arXiv*, 2020, arXiv:2008.01472.
- [10]. M. Khaneghaei, D. Asadi, Ö. Tutsoy, Software-in-the-loop simulation for an autonomous multicopter flight planning and landing with ROS and Gazebo, in *Proceedings of the International Symposium on Autonomous Systems (ISAS'21)*, 2021, pp. 1-6.
- [11]. S. M. Rahman, M. A. Hossain, H. Kim, Software-in-the-loop simulation for autonomous UAV mission validation using PX4 and Gazebo, in *Proceedings of the International Conference on Unmanned Aircraft Systems (ICUAS'21)*, 2021, pp. 142-148.
- [12]. A. Sirsavkar, M. Shaikh, A. Kale, A. Thore, et al., Autonomous UAVs for disaster response: Terrain-aware target detection and payload delivery, in *Proceedings of the International Conference on Drones and Autonomous Systems (DAUS'26)*, 2026, pp. 163-168.
- [13]. Y. Li, Q. Fan, H. Huang, Z. Han, et al., A modified YOLOv8 detection network for UAV aerial image recognition, *Drones*, Vol. 7, Issue 5, 2023, 304.



Published by International Frequency Sensor Association (IFSA) Publishing, S. L., 2026
(<https://www.sensorsportal.com>).

Universal Sensors and Transducers Interface (USTI-EXT) for extended temperature range

-55 °C ... +150 °C



26 measuring modes for all frequency-time parameters, rotational speed, capacitance Cx, resistance Rx, resistive bridges
Frequency range, 0.05 Hz ... 7.5 MHz (120 MHz);
Programmable relative error, % 1 ... 0.0005 %
Conversion speeds 6.25 us ... 12.5 ms
SPI, I2C, RS232 (master and slave, up to 76 800 baud rate)
Packages: 32-lead, 7x7 mm TQFP and 32-pad, 5x5 mm (QFN/MLF)

Applications: automotive industry, avionics, military, etc.

<http://excelera.io/> info@excelera.io

Towards Offline Navigation for Small UAVs Using Lightweight CNNs and Satellite Maps

1, 2, * Ludwig KEMPF, ¹ Mark UMANSKY, ¹ Kevin B. KOCHERSBERGER

¹ Virginia Tech, Uncrewed Systems Laboratory, 160 Inventive Lane, Blacksburg, VA 24061, USA

² Technical University of Darmstadt, Institute for Flight Systems and Automatic Control,
Otto-Berndt Strasse 2, 64287 Darmstadt, Germany

* E-mail: lkempf@vt.edu

Received: 30 Jan. 2026 /Revised: 23 April 2026 /Accepted: 25 April 2026 /Published: 28 April 2026

Abstract: Today there is a wide range of different navigation techniques with additional technologies paving the way for more complex fusion approaches. As the current standard, a combination of Inertial Navigation System (INS) and Global Navigation Satellite System (GNSS) enhanced with state-of-the-art augmentation services provide high accuracy, global coverage, and an acceptable baseline robustness. However, its reliance on signal integrity leaves safety-critical applications vulnerable to spoofing and jamming, while global crises increase the likelihood of such destructive interferences. With self-containment and high deployment volumes in mind, a lightweight vision-based reference module is proposed as a GNSS fallback for seamless integration into small Uncrewed Aerial Vehicles (UAVs). Using offline satellite maps and an ensemble of Convolutional Neural Network (CNN) models, 2D offsets and associated uncertainties are fused through the PX4 autopilot's on-board Extended Kalman Filter (EKF). Software-In-The-Loop (SITL) simulations provide first results on the feasibility and real-time capabilities of the module on a resource-limited platform. Systematic mission parameter variation on terrain data of Virginia Tech's Kentland Farm demonstrates the required minimum flight path stabilization capabilities with an R95 accuracy of 4.5 m to 5.5 m. Even in low-texture environments drift stays bounded, while varying signal loss scenarios reveal current stability limitations. Driven by the integration concept and initial SITL results, the resulting hardware test platform is presented as the foundation for upcoming flight testing.

Keywords: GNSS-denied navigation, Vision-based sensing, Cross-view geolocalization.

1. Introduction

While navigation accuracy used to be a question of technological feasibility, today it is a decision taken based on the performance demands of specific mission objectives. With standardized technology, such as the Global Positioning System (GPS) in conjunction with the Wide Area Augmentation System (WAAS), an R95 accuracy of 1.6 m can easily be achieved compared to over 20 m for a non-augmented best-case scenario with initial GPS deployment in 1995 [1, 2].

Fig. 1 illustrates the R95 metric and puts it into perspective compared to alternatives.




Fig. 1. Common accuracy measures in navigation.

With steady advances in navigation technology – from celestial navigation, to radio-based systems, to the latest intelligent approaches – fusion has always been employed in position estimation. Fig. 2 provides an overview of navigation techniques based on a literature review of methods using passive and active

perception modalities. The broadest distinction between navigation techniques can be made by absolute and relative positioning, which are usually combined with each other for complementation, as is the case for a Global Navigation Satellite System (GNSS) with an Inertial Navigation System (INS).

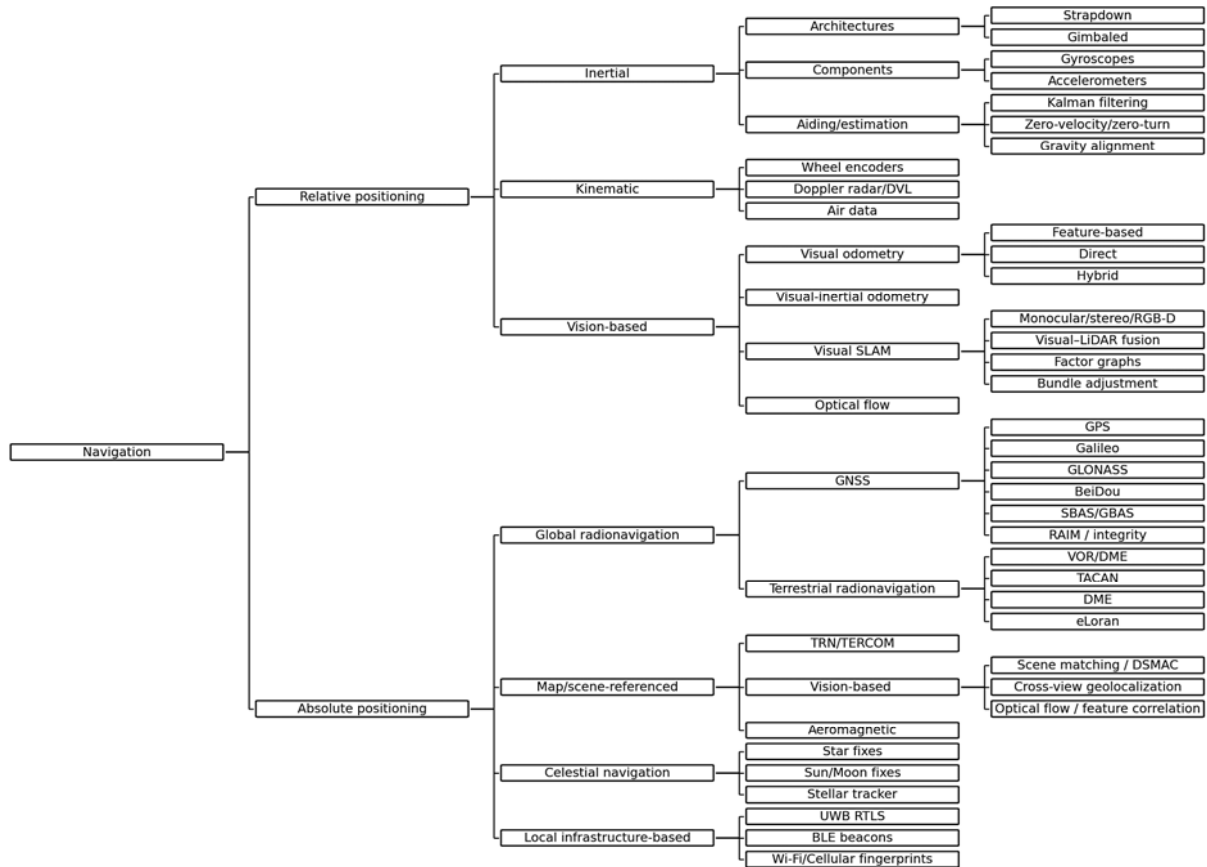


Fig. 2. Overview of navigation techniques.

Despite being self-contained, non-jammable, and non-radiating, standard INS outputs such as those of any stand-alone dead-reckoning system for relative positioning will drift over time due to bias and noise in the sensed signals and their processing. This holds true for high-end sensors and low-end sensors alike, with Fiber-Optic Gyroscopes (FOGs) representing the standard for what is possible in terms of upper price and performance, but they will still drift and require correction information provided by a global position estimator. Cost differences between low- and high-quality systems can reach up to five orders of magnitude, illustrating the economic considerations for high-volume deployment plans [3].

Building upon advanced INS technology, vision-based approaches have augmented and even replaced some conventional systems. Alongside probabilistic techniques, such as Kalman or particle filters, optical and feature-based techniques precede the most recent Artificial Intelligence (AI)-based relative positioning concepts. Nevertheless, even with today's variety of sensors available – such as cameras,

barometers, radar, or sonar – each of these imposes its own limitations [4].

In contrast, the spectrum of absolute positioning techniques stretches from the very first celestial navigation approaches to the technological extension of this triangulation concept to global radionavigation, with satellites following the first ground stations and both complementing each other in today's navigation strategies.

However, GNSS depends entirely on signal quality, and Fig. 3 illustrates the increased level of GPS (and other GNSSs) interference today compared to early March 2022. Therefore, with spoofing and jamming becoming major concerns for safety-critical applications, the need for a self-contained GNSS substitute for absolute positioning has grown and scene-referenced navigation has been identified as the most promising class of techniques. Vision-based approaches with lower hardware cost are suitable as a lightweight GNSS alternative for absolute positioning of small Uncrewed Aerial Vehicles (UAVs).

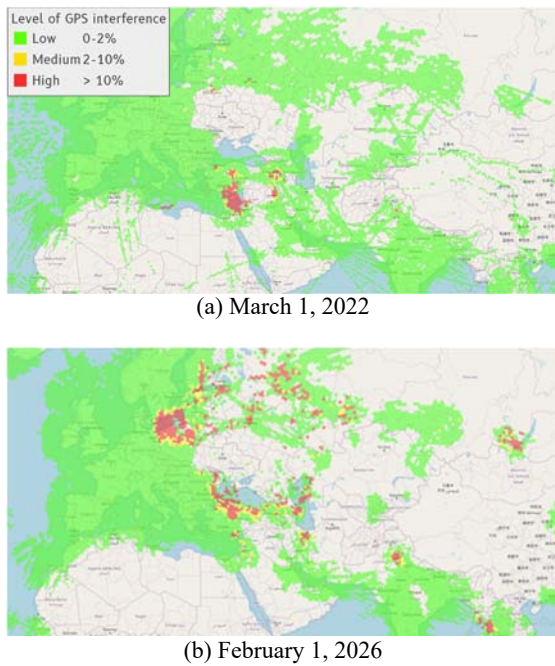


Fig. 3. Level of GPS (and other GNSSs) interference estimated based on the Navigation Integrity Category (NIC) subscribed to via Automated Dependent Surveillance Broadcast (ADS-B) Exchange API.

1.1. Scene-Referenced Navigation

The very first Terrain Referenced Navigation (TRN) systems, such as Terrain Contour Matching (TERCOM) and Digital Scene Matching Area Correlation (DSMAC), were developed during the Cold War and have been used in GNSS-denied environments ever since, albeit with varying relevance. While TERCOM compares an offline Digital Elevation Model (DEM) to actual radar altimeter readings and is usually used for en-route navigation, DSMAC exploits correlation information of binary images. Consequently, the latter excels on terrain with less variation in ground elevation but structured textures present and is favored for final approach scenarios. Typically, R95 accuracy is approximately 3 times higher for DSMAC [5, 6]. However, reliable operational specifications for TERCOM and DSMAC are not publicly available and reported accuracy values should therefore be treated as indicative estimates rather than strict performance guarantees.

Aeromagnetic navigation is formally not classified as TRN but it is still a field-referenced method with increasing relevance for global self-contained operation due to recent advances in Quantum Scalar Magnetometers (QSMs). These are required to identify minor anomalies in Earth's magnetic field masked by any aircraft's electromagnetic noise. However, when expressed in terms of R95 accuracy, the corresponding radii are more than two orders of magnitude greater than those for DSMAC or legacy single-frequency GPS without any augmentation [7].

Substantial advances in Artificial Intelligence (AI), Machine Learning (ML), and in particular Neural Networks (NNs), have enabled a variety of even more powerful scene-referenced approaches, with the two predominant categories being Template and Feature Matching (TFM) and Semantic Mapping and Recognition (SMR). As a high-level cognitive process, the latter targets content rather than structure alone, exploiting the available data to a greater extent.

As a TFM approach, cross-referencing processed RGB images with a DEM using Simultaneous Localization and Mapping (SLAM) with bundle adjustment on sufficiently high image quality, a converted R95 accuracy of less than 1 m is achievable [8]. Alternatively, using Synthetic Aperture Radar (SAR) images in combination with the range-Doppler algorithm, a converted R95 accuracy of approximately 2.5 m has been demonstrated under ideal conditions, provided that high-resolution SAR reference maps (or known reflectors) are available [9]. Fusing precise TRN, a radar altimeter and LiDAR, converted R95 accuracies between 6.4 m and 17.5 m have been measured for a variety of mission parameters [10].

In the SMR domain, using only a downward-facing monocular camera, an altimeter and a compass to compare measurements against open-source vector maps, a converted R95 accuracy of 10 m has been reported [11]. Employing object-level geolocalization, point-matching Convolutional Neural Networks (CNNs), and pixel-level refinement, a converted R95 accuracy of 0.2 m is achievable under highly favorable conditions, noting that performance depends on the evaluation dataset [12].

Table 1 compares a selection of relevant approaches. Note that, while the accuracy values have been estimated based on reasonable assumptions, they do not correspond to the metrics in which performance was originally reported. More importantly, none of the techniques have been evaluated under identical conditions (and rarely even on the same datasets).

Table 1. Converted R95 accuracies assuming an isotropic Gaussian error with Rayleigh distribution.

Method	R95 accuracy in m
Others [8, 9, 10, 11, 12]	0.2 - 30
GNSS [1, 2]	1.6 - 36
DSMAC [6]	≤ 20
TERCOM [5]	~ 60
Magnetic anomalies [7]	200 - 800

1.2. Design Approach

Guided by the introduced limitations, the proposed concept promotes an easily integrable and deployable module with a high absolute positioning accuracy threshold to bound drift, uncertainty estimation for stable Extended Kalman Filter (EKF) fusion, and real-time capability on a small UAV platform with limited computational resources. The system scales to

different mission scenarios, maintains reasonable robustness, and uses simple, inexpensive sensors, as well as satellite imagery of only moderate resolution.

With nearly 11000 small commercial satellites (≤ 1200 kg) launched between 2015 and 2024, not only has communication technology advanced, but the commercial availability of high-resolution satellite imagery has also increased significantly [13]. Planet Labs PBC alone launched more than 500 satellites during this period, and its SkySat constellation of 15 units can revisit almost any location up to ten times per day [14]. Its Dove fleet provides near-daily coverage of the entire globe [15].

Only through these recent trends providing multiple image updates of the same location every day can the introduced referencing method become a realized concept for future navigation strategies. Therefore, success ultimately depends on providers like Planet Labs PBC or BlackSky (another smaller yet highly relevant provider) supplying new data at high frequencies and reasonable costs, with current trends being very promising [16, 17].

2. Methodology & Modeling

The reference module is designed to directly integrate with the widely accepted and highly compatible PX4 open-source autopilot and its built-in EKF2, which exchanges information with the reference module via ROS2. Fig. 4 shows the methodology blueprint of the reference module. With an actively gimbaled, north-aligned RGB camera providing an observation, the observed scene is then compared to a window of the same extent as the camera frame extracted from a reference satellite map based on the last EKF position estimate. Next, both images are preprocessed and passed to the estimator ensemble, which outputs a mean 2D offset $\mu_{\Delta x}$ alongside its covariance $\Sigma_{\Delta x}$. Using the built-in `VehicleOdometry` topic, the EKF2 on the PX4 fuses the probabilistic estimate and ensures correct allocation in time using the reported timestamps.




Fig. 4. Methodology blueprint.

2.1. Design Requirements

The module is designed to provide absolute position reference not only when flying across geometrical human-made structures, but also when

flying across organic textures. However, any extreme alteration of the terrain or substantially compromised visibility (bright reflections, dark shadows, changed vegetation, clouds, fog, smoke) will put mission success at risk. The employed model ensemble is therefore specialized to the defined mission settings – in particular the type of terrain – and a broader validity range can only be established through ensemble pooling. Since the system is equipped only with an RGB camera, nighttime operations are also currently not supported. Nevertheless, other ranges and resolutions of the electromagnetic spectrum alongside additional filtering techniques are to be explored, potentially leading to extended capabilities in the future.

Since the cruise segment usually dominates the mission profile, it is the most relevant phase for navigation problems. Therefore, the initial reference module explicitly excludes any radical UAV maneuvers and targets a constant altitude Above Ground Level (AGL) of 120 m, providing sufficient spatial context for absolute positioning. The only change in flight conditions that will be further investigated within the scope of this work are wind gusts affecting gimbal stability. A conservative propagation percentage of 50 % is assumed for simulation leading to slight image out-of-plane distortions from changes in pitch and roll that are modeled using a small-angle approximation. These changes are typically below 5° , with occasional rotations up to 7° . All other effects are neglected. This approach ensures a minimum level of robustness under nominal flight conditions, but more variation during training can later enable true adaptability to greater AGL fluctuations or additional effects from a violated cruise flight condition.

Another aspect of the reference module to be closely monitored is the data used for training, currently limited to 9 satellite image sources during this proof-of-concept phase. The chosen sources and the ones thereby excluded determine the strengths and limitations of the employed models. Therefore, it is crucial to acknowledge the inherent bias from data availability and understand its effects on performance with varying terrain data during testing. Additionally, constraints such as overcast weather or satellite constellation limitations during data acquisition are two very immediate factors to be aware of.

Initial testing is conducted in Software-In-The-Loop (SITL) simulation before moving to the hardware stack for flight testing under real-world conditions. These first results provide the baseline for the following compatibility aware hardware adjustments but should at the same time be treated as highly preliminary – in particular as regards latency evaluation. EKF2 accepts signals as old as 1000 ms, allowing for seamless SITL testing without optimizing, but a translation from Python, used for fast prototyping, to another more efficient language, such as C++, may be necessary to meet the 10 Hz target data rate for final deployment.

2.2. Data Sampling

A total of six satellite images of Virginia Tech’s Kentland Farm were selected based on its visual correspondence to the test mission from a pool of 9 images (Fig. 5) provided by BlackSky [18]. While the reference satellite map is selected to reflect the latest changes in terrain (1), additional sources are identified with training convergence in mind but still allowing for anticipated fluctuation in lighting conditions (2-4). At the same time, robustness is maintained by also including more significant vegetation and foliage changes (5, 6).

Overall, this yields a 200000-sample dataset of paired satellite windows and surrogate camera frames, downsampled from 1 m to 3.8 m per pixel for low-pass filtering and converted to greyscale. Within the defined sampling domain – enabling a later cross-area evaluation – window centers and associated offsets are randomly chosen. Offsets are bound to ± 10 m (8.3 % of the frame size) on both Cartesian axes to optimally use the limited complexity of the CNN.

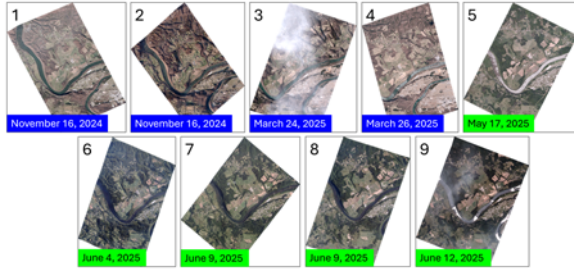


Fig. 5. Raw satellite images sorted by date. Blue labels show low vegetation coverage, and green labels show high vegetation coverage corresponding to seasonal change.

2.3. Gaussian Model

The individual models are trained following a standard supervised curriculum learning approach, with the first 10 epochs training only the means using the mean model, and the next 30 epochs extending the output to uncertainties using the Gaussian model and a custom Negative Log Likelihood (NLL) loss. Fig. 6 shows the overall high-level architecture.

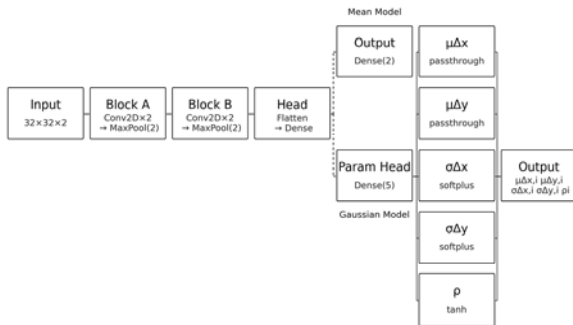


Fig. 6. CNN architecture.

The architecture’s overall simplicity allows a straightforward interpretation of performance, representing a solid baseline for the future implementation of more advanced structures and optimization of estimate quality.

An additional calibration term enforces the average squared Mahalanobis distance m^2 to remain close to 2 with the objective of well-scaled uncertainty coverage and Fig. 7 shows the Cumulative Distribution Function (CDF) of m^2 , representing the probability of the value being less than or equal to a given threshold. The empirical curve computed from the validation data generally follows the theoretical distribution, indicating that the estimated uncertainties have the correct order of magnitude. Generalization of calibration must later be assessed on the test data.




Fig. 7. Calibration curve of the Mahalanobis distance squared m^2 of a typical training run.

2.4. Estimator Ensemble

To constitute the estimator ensemble, 3 separate CNN models are trained on the same dataset but with different sample shuffling, so that the variance across their predictions $\Sigma_{\mu_{\Delta x}}$ provide an inexpensive estimate of systematic uncertainty from model disagreement. This component can then be combined with the average of the stochastic uncertainty $\bar{\Sigma}_{\Delta x}$ of all models M to form the total effective covariance:

$$\Sigma_{eff,\Delta x} = \bar{\Sigma}_{\Delta x} + \Sigma_{\mu_{\Delta x}}, \quad (1)$$

with

$$\bar{\Sigma}_{\Delta x} = \frac{1}{M} \sum_{i=1}^M \Sigma_{\Delta x,i}, \quad (2)$$

and

$$\Sigma_{\mu_{\Delta x}} = \frac{1}{M} \sum_{i=1}^M (\mu_{\Delta x,i} - \bar{\mu}_{\Delta x})(\mu_{\Delta x,i} - \bar{\mu}_{\Delta x})^T \quad (3)$$

The diagonal elements of the covariance ultimately guide EKF updates and signal fusion on the PX4.

Although systematic uncertainty is already modeled, there remains the risk that all models in an ensemble may share the same bias. This results in an unreasonable estimate that, in turn, produces large innovations within the EKF2. To safeguard against such scenarios and maintain filter consistency, an innovation-based covariance inflation is applied, with the normalized innovation test ratios τ_x and τ_y fed back into the reference module. The ratios per axis are calculated from the innovation v , describing the disagreement between EKF motion prediction and the estimator ensemble output, and the innovation gate γ . The latter is an integer defining the acceptable multiples of σ (eq. (4)) and is set to 5 – the PX4 default value.

$$\tau_i = \frac{v_i^2}{\gamma^2 \sigma_i^2} \quad (4)$$

When these ratios become excessively high, $\Sigma_{eff, \Delta x}$ is inflated by a factor λ , as defined in eq. (5). This adaptive mechanism enhances numerical stability and preserves statistical consistency. Ideally, these ratios remain below 0.5. Values exceeding 1 indicate potential overconfidence of the estimator and lead to the rejection of the measurement.

$$\lambda = 1 + \max(0, \tau_{max} - 0.5), \quad (5)$$

with

$$\tau_{max} = \max(\tau_x, \tau_y) \quad (6)$$

This logic has an intuitive implication: Within a single update step, substantial disagreement is tolerated only when uncertainty is high. Nevertheless, it does not prevent destabilization when motion prediction and estimator output are close to each other but both wrong.

The flowchart in Fig. 8 summarizes the estimator ensemble subprocess, outlining the most important steps.




Fig. 8. Flowchart of the estimator ensemble subprocess.

3. Simulation Results

For the systematic evaluation of the reference module, real-time SITL simulation with the X500 Quadrotor model in Gazebo Harmonic is performed on the Sseed Technology reComputer J4012, which is based on the NVIDIA Jetson Orin NX 16 GB module [19]. A dedicated geo-mosaic of Kentland Farm together with its corresponding DEM for relief modeling (virtual object displacement) constitutes the terrain data for simulation. This allows for a simulated camera view for a variety of mission scenarios (Fig. 9). The raw data was collected using a SenseFly eBee fixed-wing mapping drone.




Fig. 9. Geo-mosaic and associated DEM for camera simulation during SITL testing.

SITL testing starts with an open-loop reference mission, assessing the reference module's performance compared to ground truth but still using GNSS for positioning. The reference is then modified step by step to explore the following modifications:

1. Loop closure (GNSS disabled);
2. Mission location (cross-area);
3. Linear flight path geometry;
4. Wind gust exposure;
5. Signal loss:
 - a. Linear w/ wind;
 - b. Linear w/o wind;
 - c. Circular w/o wind.

This process identifies relevant trends and examines robustness of the proposed reference module by better understanding its strength and limitations. Fig. 10 shows the in-area reference mission, with at least 75 % of the terrain covered by pastures and farmland. The Ground Control Points (GCPs) indicate the centers of the 200 m sampling radii for the employed dataset. As for all subsequent variations, the ground speed is set at 5 m/s.

The statistics in Fig. 11a demonstrate a cumulative R95 accuracy of 4.7 m along the course of the reference mission, including the spike about 400 s into the flight. Given the final resolution of 3.8 m, the mean bias of 1.2 m stays in the sub-pixel range. At the same time, the coverage statistics reveal the reference module's constant overconfidence on unseen data. There is a significant under-coverage compared to the expected percentages (hereafter denoted exp) for all three bands and Fig. 11b shows the sudden disagreement between the EKF2's prediction and the external position estimate after a series of lower

innovation test ratios. However, test ratios stay within a healthy regime throughout the mission.




Fig. 10. Route of the reference mission. GCPs indicate the centers of the 200 m sampling radii for the employed training dataset.


3.1. Loop Closure

When disabling GNSS and thereby closing the control loop, R95 accuracy declines, with a value of 5.1 m for the selected run. This represents an average scenario, but results vary by approximately 0.5 m across different runs.

As anticipated, being a function of the ground textures, coverage is almost identical to the reference mission. Although the noisy estimator outputs induce frequent deviations from the target, the position estimate remains stable and within the indicated accuracy. Fig. 12 shows a typical flight-path anomaly caused by noisy estimates. The blue band indicates the standard deviation obtained by projecting the effective covariance $\Sigma_{eff,\Delta x}$ onto the unit vector perpendicular to the trajectory. It illustrates that while the estimator tracks the overall trend, the 1σ band does not provide sufficient uncertainty coverage.



(a) L2 offset of the estimated position with respect to the truth. The blue band around the estimator represents the radial standard deviation obtained by projecting the effective covariance $\Sigma_{eff,\Delta x}$ onto the offset direction unit vector



(b) Innovation test ratios

Fig. 11. Time metrics for the reference mission.

Computing the relevant ground truth statistics, Fig. 13a uncovers the cause of drift at GCP 12 (400 s into the flight) for the reference and the loop-closure mission: there is a significant bias in the negative x-direction from training-to-truth inconsistency.

Although Fig. 11a earlier revealed significant under-coverage, Fig. 13b now shows that the squared Mahalanobis distance m^2 has the correct order of magnitude in most regions, which is a very promising tendency in the data.




Fig. 12. Flight path anomalies for the loop-closure mission.

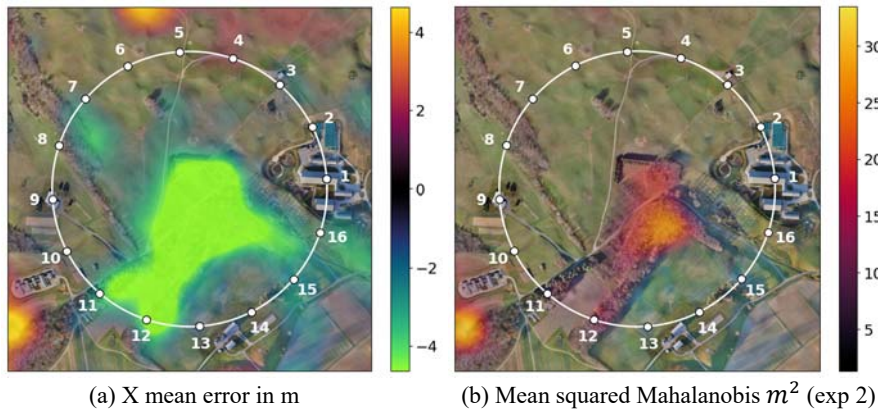


Fig. 13. Ground truth statistics for the reference mission, generated by performing repeated Monte Carlo simulations of the estimator ensemble on discrete grid cells across the geo-mosaic. Grid resolution is 60×60 m, with 30 runs per cell and a maximum offset of 5 m for regularization.

Fig. 14 illustrates the UAV's behavior at the alternate mission location. After a stable flight, only showing the typical anomalies, a significant and persistent disagreement between the truth and the estimate begins to manifest up to a destabilization.

With the L2 offset Fig. 15 illustrates the gradual unfolding of the destabilization in the time domain.

The 1σ -band indicates a generally larger reported uncertainty over the forest region, resulting in lower innovation test ratios and smaller updates of the motion-predicted position. Consequently, the EKF2 leans more heavily on these predictions and becomes overconfident, with the destructive divergence appearing around 370 s into the flight. During this phase, the motion model and estimator ensemble output agree with each other while both being wrong.

While the models still perform well – even in previously unseen areas – if similarity is provided, a significant inconsistency for an extended period will break the EKF logic and therefore induce a high risk of destabilization (Fig. 16).

3.3. Linear Flight Path

Changing the flight path geometry of the mission could be nontrivial, since a straight line requires less control effort than a continuous circle. However, with

3.2. Mission Location

While all simulation is conducted as an out-of-sample test on unseen data, the alternate location additionally represents a cross-area evaluation on a region that contains mostly sections the employed models have not been trained on. In contrast to the reference mission, this mission leads across forests and has a tree coverage of 20 %, with the most challenging part being the continuous nature of the relevant passage. GNSS is disabled (as for all missions except for the reference mission).

the PX4 already being a highly optimized autopilot with a sophisticated cascade control architecture, the flight path geometry doesn't affect the performance of the reference module for a continuous position signal [20]. Later investigation shows that it becomes relevant for a signal loss scenario.




Fig. 14. Destabilization across forest.

3.4. Wind Gust Exposure

For a robust system, wind should only have a minor influence on its functionality. Wind gusts

represent one of the most challenging scenarios as they can have a significant impact on the flight state – particularly on the pitch and roll of the quadrotor

model. These dynamic attitude changes must not compromise the output quality of the estimator ensemble under nominal flight conditions.




Fig. 15. L2 offset of the estimator relative to the truth for the alternate mission location.

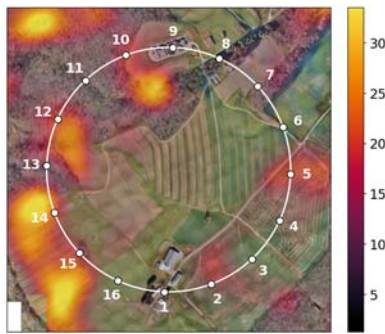


Fig. 16. Mean squared Mahalanobis m^2 for the alternate mission location (exp 2).

Despite the actively controlled gimbal being part of the system design, strong gusts may still impose a real risk. Therefore, a conservative propagation percentage of 50 % is assumed for simulation and slight image distortions from changes in pitch and roll are modeled using a small-angle approximation. These changes are typically below 5° , with occasional rotations up to 7° . Yaw is assumed constant.


Gazebo Harmonic has limited aerodynamic modeling capabilities compared to Classic (end of life in 2025), so the included `windEffects` plugin is empirically tuned in a way that simulates a realistic wind gust scenario for a small UAV. Starting from a 7 m/s constant wind baseline, strong gusts are induced.

It shows that high-frequency fluctuations of the position observation caused by wind gusts are absorbed by the noise of the estimator ensemble itself and therefore do not compromise the reference module's estimate quality.

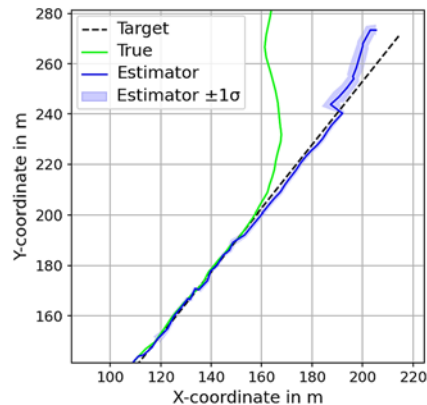
3.5. Signal Loss

Undetected signal loss or skipped measurements over extended periods pose potentially catastrophic

consequences for flight stability and mission success. Although PX4 will initiate an auto-land once a missing external position estimate is detected, it will continue flying if the signal appears valid. This scenario is simulated with a 10 s blackout and a corresponding estimate freeze. Fig. 17 shows the behavior without wind and with additional wind gust exposure.



(a) Recovery without wind




(b) Destabilization with wind


Fig. 17. Behavior during a 10 s signal loss Initiated around $x = 150$ m.

While the system generally destabilizes under wind-gusts conditions, even on a linear flight path, it has a high likelihood of recovering when no gusts are

present. On a circular flight path without wind, the UAV is less likely to recover after a signal loss, although it is possible as captured in Fig. 18.



(a) L2 offset of the estimator relative to the truth



(b) Innovation test ratios

Fig. 18. Time metrics for signal loss after 50 s on the circular path without wind.

The notable spike in the y innovation test ratio indicates the estimator ensemble having regained a valid input signal. However, the EKF still rejects it due to its magnitude. Once innovation-based covariance inflation takes effect, the filter accepts the estimator outputs again, and PX4 regains control.

4. Future Test Design

Overall R95 accuracy for the different test cases ranges from 4.5 m to 5.5 m, with flights across distinct human-made structures typically improving accuracy by about 50 %. At the same time, ground truth evaluation suggests that despite performance variation by terrain, the estimator ensemble's baseline uncertainty has the correct order of magnitude to provide long-term EKF consistency.

While additional ablation studies are still to be conducted, it shows that using an ensemble instead of a single model is a promising approach, especially for distinct textures (Fig. 19). However, extended forest coverage imposes a high risk of destabilization without further measures taken. Prior automated

terrain classification and specialized CNN models could improve coverage in more complex regions.

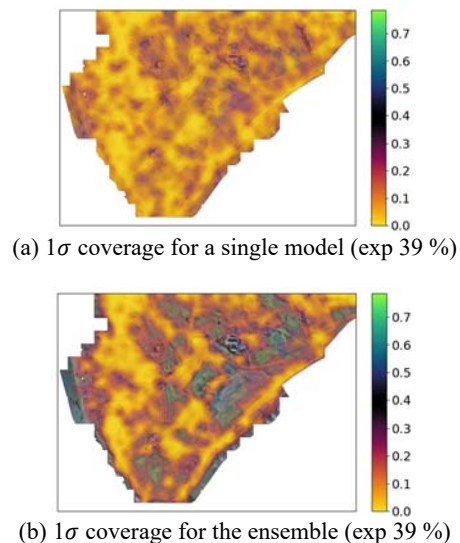


Fig. 19. Coverage overview for a single model and the estimator ensemble.

For future testing, instead of looking at a realistic mission scenario with heterogeneous terrain, individual ground textures should be evaluated by simulating test missions on homogenous terrain, with each simulation using a model ensemble solely trained on representative samples. Together with stronger out-of-distribution testing, this would allow for a systematic performance evaluation by terrain, therefore providing the foundation for purposefully increased model and model ensemble complexity.

Additionally, the reference module's current design does not provide any global search – as opposed to GNSS allowing full localization at any point in time. With the proposed basic CNN architecture, results are easier to interpret but the module has a very limited receptive field of ± 10 m, providing a drift-bounding reference for absolute positioning only when a rough initial guess is available for reference. Therefore, future test design should evaluate a cascade approach with sequential focusing of the receptive field, thereby refining the search. Visual Transformers (ViTs) could be used for an upstream global search instead of traditional CNNs.

Although closed-loop flight tests had not been performed at the time of article submission, the Tarot T960 hexacopter platform was selected as the base for integrating the rest of the hardware system for Hardware-in-the-Loop (HITL) simulation and future flight tests, as shown in Fig. 20. It provides a payload capacity of about 4 kg and flexible mounting options on the frame are available for adding additional devices such as computers, radios, and sensors.



Fig. 20. The integrated Tarot T960 with camera payload and computing platform.

The Tarot's six brushless motors are controlled by a PixHawk Cube Black (Fig. 21) running the latest stable PX4 version 1.16.0, the exact same version used for SITL evaluation. It connects to the Electronic Speed Controllers (ESCs) for commanding the motors based on its control inputs or preplanned autonomous mission.

With an observed loop rate of 5.7 Hz on average (min 3 Hz, max 10 Hz), the target data rate of 10 Hz is not currently achievable. Since model inference alone takes about 25 ms per cycle, a transition to the Graphics Processing Unit (GPU) and from Python to embedded C++ cannot be the only modification. The NVIDIA Jetson Orin AGX 32GB is selected to replace the NVIDIA Jetson Orin NX 16 GB as an enduring main computing platform. The Orin AGX consumes

up to 40W while delivering up to 200 TOPS of processing power, compared to 157 TOPS available in the Orin NX 16GB [29]. The Orin AGX is connected to an Auvidia X230 carrier board, providing the required mechanical, electrical, and communications interfaces for integrating into the Tarot (Fig. 22).



Fig. 21. PixHawk Cube Black [21].



Fig. 22. Auvidia X230 carrier board with Orin AGX [22].

The carrier board connects to the PixHawk Cube Black over a standard serial connection and to the Tarot's main camera sensor over gigabit ethernet. An additional code and language optimization might still be necessary to minimize the overhead from statistical computations of the model ensemble.

Instead of just using an RGB camera, Infrared (IR) could be used for nighttime missions and polarization information could be exploited to improve vision in foggy scenes [23, 24]. The latter will be implemented in future testing, with the primary sensing input being a LUCID Vision Phoenix PHX050S1-QC polarimetric camera. The camera contains a 5MP Sony IMX264MYR CMOS sensor, which is a color sensor that has an array of micro-lenses with linearly polarizing filters on each 4-pixel super-pixel (Fig. 23). It is powered either over its gigabit ethernet connector via PoE or directly via DC voltage.

A custom camera driver has been written using the LUCID Vision Software Development Kit (SDK) to connect to the camera, configure its acquisition parameters, and produce a standard ROS2 image interface.

While SITL simulation shows that gust-induced fluctuations in pitch and roll of less than 7° are tolerated (max 3 % pixel shift), in-plane camera variation from yaw imposes a real challenge. The ensemble is not trained on off-axis distortions. Therefore, to provide a stable observation, the camera is mounted to a 3-DOF HDAir Studio MR-S2 gimbal keeping the sensor north-aligned and pointed downwards in the nadir direction during typical flight maneuvers. The gimbal is preconfigured which removes any need for active external command and control of the gimbal.




Fig. 23. Sony IMX278 Polarization Sensor [25].

Alternatively, an orthorectification algorithm could be a cost-effective gimbal replacement for a future flight stack and has been proven feasible in past research [26]. Nevertheless, to ensure robustness against long-term magnetometer inaccuracies, an additional yaw Degree Of Freedom (DOF) must be included in the training of the models themselves.

Like in SITL simulation, all communication between the PX4 and any external nodes is established through the micro Unified eXtensible Resource Coordination Entity Data Distribution Service (uXRCE-DDS) as illustrated in Fig. 24.




Fig. 24. PX4/ROS2 link over uXRCE-DDS [27].

5. Conclusion & Outlook

While dynamically changing terrain such as tree covered regions is identified as a major challenge, the proposed reference module can still effectively bound drift when flying across low-texture farmland and pastures. The overall R95 accuracy stays between 4.5 m and 5.5 m and the squared Mahalanobis distance has the correct order of magnitude for most of the test terrains. Features persisting across training and truth are the fundamental precondition. However, a systematic classification of terrain-performance relationship is required to push the reference module towards its full potential and boost robustness of uncertainty coverage on unseen data. Additionally, only an upstream global search would transform the current module to a full GNSS fallback. Constrained computational efficiency and sensitivity towards yaw

must be addressed as well. Comprehensive benchmarking against alternative techniques, additional ablation studies, and larger datasets for training and testing are required to further substantiate the results and identify potential improvements.

Initial integration tests of the Tarot T960 base flight platform, Orin computing system, camera sensor and gimbal, camera driver, and PX4 communications have been completed, with short test hop flights successfully executed to verify the overall stability of the mechanical and electrical integration of the payloads. The next integration step is to close the loop with the reference module providing absolute position data in real time to the PX4's EKF2. This will enable iterative experiments with actual flight and sensor data, as well as evaluation and extension of the results from SITL simulations completed thus far.

Acknowledgements

The presented research is based on the first author's thesis submitted in partial fulfillment of the requirements for the degree of Master of Science in Mechanical Engineering at Virginia Tech (USA) and Technical University of Darmstadt (Germany). The thesis is currently under embargo.

Fundamental findings have already been presented at the *2nd International Conference on Drones and Unmanned Systems (DAUS' 2026)*, with its associated conference proceedings [28]. This article extends the proceedings as follows:

1. Additional background on state-of-the-art navigation techniques and recent motivation is provided;
2. The integration of adaptive covariance inflation is explained and the full model ensemble process illustrated in a flow chart;
3. All scenarios simulated so far are discussed, now including a GNSS-enabled reference mission, a variation in flight path geometry, and wind-gust scenarios with and without signal loss;
4. Future steps in SITL simulation are explicitly derived based on the limitations revealed from the initial results;
5. A specific hardware setup has been developed to provide a test platform for HITL simulation and flight tests in the field.

The research is funded by the National Science Foundation Center for Autonomous Air Mobility and Sensing (CAAMS) and supported by a fellowship of the German Academic Exchange Service (DAAD).

References

- [1]. Global Positioning System precise positioning service performance standard, *United States Department of Defense*, 2007.

- [2]. Global Positioning System Wide Area Augmentation System (WAAS) performance standard, *Federal Aviation Administration*, 2008.
- [3]. Advanced Navigation, MEMS vs FOG: What inertial system should you choose?, <https://www.advancednavigation.com/tech-articles/mems-vs-fog-what-inertial-system-should-you-choose/>
- [4]. I. Jarraya, A. B. Mohammed, H. N. Noura, O. Salman, et al., GNSS-denied unmanned aerial vehicle navigation: Analyzing computational complexity, sensor fusion, and localization methodologies, *Satellite Navigation*, Vol. 6, Issue 1, 2025, 9.
- [5]. Worldwide equipment guide, Volume 2: Airspace and air defense systems, *Department of the Army*, 2011.
- [6]. J. C. O'Halloran (Ed.), Jane's Weapons: Strategic 2015-2016, *IHS Global Inc., Coulsdon*, 2015.
- [7]. M. Muradoglu, M. T. Johnsson, N. Wilson, T. Navickas, et al., Quantum-assured magnetic navigation achieves positioning accuracy better than a strategic-grade INS in airborne and ground-based field trials, *arXiv*, 2025, arXiv:2504.08167.
- [8]. X. Wan, M. Zhang, Y. Wang, B. Li, Terrain aided planetary UAV localization based on geo-referencing, *IEEE Transactions on Geoscience and Remote Sensing*, Vol. 60, 2022, pp. 1-18.
- [9]. C. Lindstrom, J. L. P. Barker, T. E. Humphreys, GPS-denied navigation aided by synthetic aperture radar using the range-Doppler algorithm, *NAVIGATION: Journal of the Institute of Navigation*, Vol. 69, Issue 3, 2022, 533.
- [10]. J. Lee, C.-K. Sung, J. Oh, K. Han, et al., A pragmatic approach to the design of advanced precision terrain-aided navigation for UAVs and its verification, *Remote Sensing*, Vol. 12, Issue 9, 2020, 1396.
- [11]. C. Ouyang, Y. Wang, J. Chen, X. Wang, et al., A semantic vector map-based approach for aircraft positioning in GNSS/GPS denied large-scale environment, *Defence Technology*, Vol. 34, 2024, pp. 1-10.
- [12]. Q. Ye, J. Luo, Y. Lin, A coarse-to-fine visual geo-localization method for GNSS-denied UAV with oblique-view imagery, *ISPRS Journal of Photogrammetry and Remote Sensing*, Vol. 212, 2024, pp. 306-322.
- [13]. Smallsats by the numbers 2025, BryceTech, <https://brycetech.com/reports>
- [14]. SkySat documentation, <https://docs.planet.com/data/imagery/skysat/>
- [15]. PlanetScope documentation, <https://docs.planet.com/data/imagery/planetscope/>
- [16]. Form 10-K, fiscal year ended December 31, 2024, *BlackSky Technology Inc.*, 2025.
- [17]. Form 10-K, fiscal year ended January 31, *Planet Labs PBC*, 2025.
- [18]. BlackSky imagery and analytics: Satellite images of Kentland Farm (Virginia Tech), *BlackSky Technology Inc.*, 2025.
- [19]. Seeed Technology, NVIDIA Jetson Orin NX 16GB AI device – reComputer J4012, <https://www.seeedstudio.com/reComputer-J4012-p-5586.html>
- [20]. Controller diagrams, https://docs.px4.io/main/en/flight_stack/controller_diagrams
- [21]. Hex Cube Black flight controller, https://docs.px4.io/main/en/flight_controller/pixhawk-2
- [22]. Auvideo X230, <https://auvideo.eu/product/x230-70419/>
- [23]. J. Fan, E. Zheng, Y. He, J. Yang, A cross-view geo-localization algorithm using UAV image and satellite image, *Sensors*, Vol. 24, Issue 12, 2024, 3719.
- [24]. M. Garcia, T. Davis, S. Blair, N. Cui, et al., Bioinspired polarization imager with high dynamic range, *Optica*, Vol. 5, Issue 10, 2018, pp. 1240-1246.
- [25]. LUCID Vision Labs, Polarization explained, <https://thinklucid.com/tech-briefs/polarization-explained-sony-polarized-sensor/>
- [26]. J. Kinnari, R. Renzulli, F. Verdoja, V. Kyrki, LSVL: Large-scale season-invariant visual localization for UAVs, *Robotics and Autonomous Systems*, Vol. 168, 2023, 104497.
- [27]. PX4 ROS 2 user guide, https://docs.px4.io/main/en/ros2/user_guide
- [28]. L. Kempf, K. B. Kochersberger, A lightweight vision-based GNSS substitute for small UAVs, in *Proceedings of the 2nd International Conference on Drones and Unmanned Systems (DAUS'26)*, 2026, pp. 173-178.
- [29]. NVIDIA Jetson Orin, <https://www.nvidia.com/en-us/autonomous-machines/embedded-systems/jetson-orin/>



Joint Rotor Speed Scheduling and Trajectory Optimization for Energy-Efficient Electric UAV Rotorcraft Missions

* Emeka CHIJOKE and Marcin ŻUGAJ

Warsaw University of Technology, pl. Politechniki 1, 00-661 Warsaw, Poland
Tel.: +48 222347220

* E-mail: emeka.chijioke.dokt@pw.edu.pl

Received: 19 Jan. 2026 /Revised:13 April 2026 /Accepted: 20 April 2026 /Published:28 April 2026

Abstract: Electric unmanned rotorcraft are increasingly considered for aerial mobility applications, but their performance is limited by onboard energy storage, making energy efficiency a critical design consideration. This study investigates the effects of rotor speed scheduling and trajectory optimization on total mission energy consumption for a hover–ascent–cruise–descent profile. A mission-based energy model is developed to estimate rotor power and cumulative energy across all flight phases. A sequential optimization framework is adopted, in which rotor speed optimization is first performed to determine a phase-dependent rotor speed schedule, followed by trajectory optimization conducted under a fixed rotor speed profile obtained from the preceding stage, with only trajectory parameters allowed to vary. A joint optimization is then evaluated to assess the integrated effect of both strategies. Results show that rotor speed optimization reduces energy by 21.85 %, while trajectory optimization achieves a larger reduction of 24.99 %. The joint optimization yields the lowest energy consumption, with a total reduction of 25.01 %, providing only a marginal improvement over trajectory optimization. These findings indicate that trajectory design is the dominant contributor to energy savings, while rotor speed scheduling offers complementary benefits for improving rotorcraft endurance.

Keywords: Rotor speed scheduling, Trajectory optimization, Energy consumption, Energy-efficient flight, Electric rotorcraft, UAV, Mission energy optimization.

1. Introduction

This article is an extended version of the conference paper presented in [1], where a preliminary version of this work was introduced, including the staged optimization framework and initial simulation results. The present work includes additional analysis, expanded methodology, extended results, and deeper discussion of the coupled effects between rotor speed scheduling and trajectory optimization.

The methodology is further refined through the formulation of a clearly defined sequential optimization framework, in which rotor speed optimization is performed prior to trajectory optimization using a fixed rotor speed schedule, with

only trajectory parameters allowed to vary. In addition, the work includes expanded parametric analyses, improved figure presentation, and a more comprehensive comparison of energy consumption across rotor speed, trajectory, and joint optimization strategies. The results are also critically reassessed to account for numerical sensitivity and modeling limitations, providing a more rigorous evaluation of the relative contributions of each optimization approach.

Unmanned aerial vehicle (UAV) electric rotorcraft are gaining increasing attention in both civil and military aerospace applications due to their potential for reduced emissions, lower operating costs, and quieter operation compared with conventional fuel

powered platforms. These advantages have made electric vertical takeoff and landing (eVTOL) vehicles and unmanned rotorcraft particularly attractive for emerging applications such as urban air mobility, infrastructure inspection, environmental monitoring, and short-range transportation [2, 3]. Despite these benefits, the operational capability of electric rotorcraft remains strongly constrained by the relatively low energy density of current battery technologies. As a result, improving mission energy efficiency has become a critical research challenge in the design and operation of electrically powered aerial vehicles [4].

The energy consumption of a rotorcraft mission is determined by a combination of aerodynamic, propulsion, and flight mechanics factors that vary across different flight phases. Hover and climb segments typically require high induced power, while cruise flight is dominated by parasite drag and forward flight aerodynamic effects [5, 6]. Consequently, the total mission energy demand depends not only on the rotor performance characteristics but also on the trajectory followed during the mission. Operational parameters such as rotor rotational speed, flight velocity, and climb or descent rates can significantly influence the power required to sustain flight. Optimizing these parameters therefore offers a practical pathway for improving rotorcraft endurance without major hardware modifications.

One approach that has received considerable attention is rotor speed scheduling. Variable rotor speed operation allows the rotor to operate closer to its optimal aerodynamic efficiency under varying flight conditions. Higher rotor speeds may be beneficial in hover or low speed flight, while lower rotor speeds during cruise can reduce profile power and mechanical losses. Previous studies have demonstrated that adaptive rotor speed control can reduce power consumption and improve endurance in both conventional helicopters and electric rotorcraft [7-9]. However, rotor speed optimization alone does not account for the influence of the mission trajectory on aerodynamic loading and overall energy demand.

Another important strategy for improving energy efficiency involves trajectory optimization. By adjusting parameters such as cruise velocity, climb rate, and descent rate, rotorcraft can operate closer to aerodynamic conditions that minimize power requirements. Efficient trajectory planning can reduce unnecessary altitude changes, avoid unfavorable environmental conditions, and improve overall mission energy performance. Previous work has shown that routing strategies that consider wind conditions and energy-efficient path planning can significantly reduce the energy required to complete UAV missions [10].

Although rotor speed scheduling and trajectory optimization have each been studied independently [11-13], the interaction between these two control strategies has received comparatively less attention. Rotor speed directly influences rotor aerodynamic efficiency, while trajectory parameters determine the

aerodynamic loading experienced during flight. Because these variables are inherently coupled, optimizing them independently may not fully exploit the potential for mission energy reduction. In particular, it remains unclear to what extent rotor speed variation continues to provide benefits once the trajectory has already been optimized.

The objective of the present study is therefore to develop a mission-based energy model for electric rotorcraft covering hover, ascent, cruise, and descent phases. Using this model, three optimization strategies are implemented and compared, namely rotor speed optimization with a baseline trajectory, trajectory optimization using the rotor speed schedule obtained from the previous stage, and joint optimization in which both rotor speed and trajectory parameters are adjusted simultaneously. The results are analyzed to quantify the relative contribution of each strategy to mission energy reduction and to identify the dominant factors influencing rotorcraft energy efficiency.

2. Literature Review

Research on rotorcraft energy efficiency has primarily focused on three key areas, namely aerodynamic modeling of rotor performance, rotor speed control strategies, and trajectory optimization for UAV missions. These aspects are closely interconnected, as energy consumption depends not only on rotor aerodynamics but also on how the vehicle is controlled and operated along its flight path. The following sections review the main developments in each of these areas.

2.1. Aerodynamic Power Modeling

Early investigations into rotorcraft power consumption focused on the development of analytical models for estimating rotor aerodynamic performance. Momentum theory and blade element methods have been widely used to estimate induced, profile, and parasite power contributions during different flight conditions [5, 6]. These models form the foundation for predicting rotorcraft energy consumption during hover and forward flight. Subsequent studies incorporated additional aerodynamic effects such as rotor inflow dynamics, blade drag characteristics, and forward flight performance, enabling more accurate estimation of rotorcraft power requirements across complete mission profiles.

2.2. Variable Rotor Speed Operation

Several studies have investigated the benefits of variable rotor speed operation. Traditional helicopters typically operate at nearly constant rotor speeds to simplify control and maintain stable aerodynamic conditions. However, advances in electric propulsion systems and motor control technologies have enabled

more flexible rotor speed management. Adjusting rotor speed according to flight conditions can reduce power consumption by allowing the rotor to operate closer to its optimal aerodynamic efficiency [7-9]. In [14], a trajectory tracking control method for a helicopter UAV based on a kinematic model and feedback linearization was proposed. The controller adjusts forward velocity to improve tracking performance while reducing energy expenditure. Such studies demonstrate that rotor speed management can play an important role in improving UAV energy efficiency.

2.3. Trajectory Optimization

Another important area of research involves trajectory optimization for minimizing UAV energy consumption. Trajectory planning methods aim to reduce unnecessary altitude variations, minimize aerodynamic drag, and exploit favorable environmental conditions. Although trajectory optimization has been extensively studied for fixed wing UAVs [15-17], these methods cannot be directly applied to rotorcraft systems because rotor power requirements depend strongly on rotor aerodynamic loading and induced velocity effects. For rotorcraft UAVs, [18] proposed a control framework based on Model Predictive Control (MPC) that operates in real-time, allowing UAVs to adapt their trajectories dynamically while minimizing altitude variations and travel distance.

More recent studies have incorporated environmental wind effects into trajectory planning. In [10], realistic urban wind fields were generated using the Parallelized Large-Eddy Simulation Model (PALM) and integrated into an enhanced A-star search algorithm for UAV trajectory optimization. The results demonstrated that exploiting wind information during path planning can significantly reduce energy consumption compared with conventional shortest path approaches. Similarly, [19] proposed a planning framework based on an energy map that evaluates energy requirements of UAV components within a three-dimensional (3D) environment and uses heuristic graph search algorithms to determine energy-efficient trajectories. In [20], energy-efficient trajectory planning for a quadrotor UAV was formulated as an optimal control problem using a power loss model that accounts for rotor speed and acceleration.

2.4. Research Gap

Despite these advances, most existing studies optimize either rotor speed scheduling or trajectory design independently. However, because rotor speed influences aerodynamic efficiency while trajectory parameters determine aerodynamic loading during flight, these two mechanisms are inherently coupled.

Understanding their interaction is therefore essential for identifying effective strategies to reduce rotorcraft mission energy consumption. To address this limitation, the present study evaluates rotor speed scheduling, trajectory optimization, and their joint implementation within a unified mission-based energy framework.

3. Methodology

3.1. Simulation and Optimization Framework

The methodology presented in this work builds upon the framework introduced in the conference paper [1], while extending it through a more comprehensive optimization formulation and a detailed analysis of the interaction between rotor speed scheduling and trajectory design.

The simulation workflow integrates MATLAB/Simulink for supervisory scheduling and command shaping with FLIGHTLAB's nonlinear six degrees of freedom (6DoF) rotorcraft model to capture high fidelity flight dynamics and aerodynamics. The combined environment enables consistent evaluation of rotorcraft performance under varying control strategies.

Mission simulations generate time-series outputs including altitude, airspeed, vertical speed, rotor speed, power consumption, and cumulative energy. These outputs are interpolated to a common time base to allow direct comparison across different optimization scenarios. An overview of the methodology workflow is shown in Fig. 1.

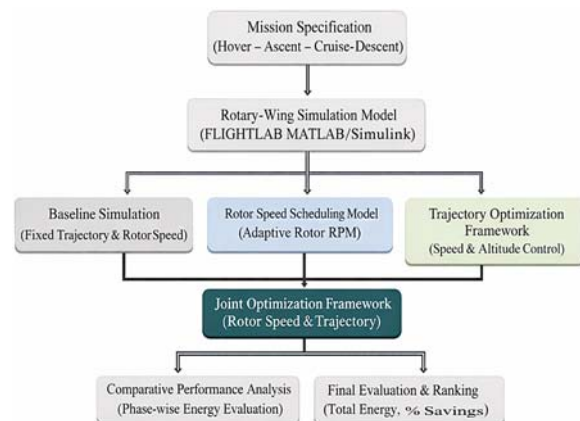


Fig. 1. Optimization framework for rotor speed and trajectory design (adapted from [1]).

3.2. Mission-Based Energy Modeling

A mission-based energy modeling approach is used to evaluate rotorcraft power consumption across different flight phases, including hover, ascent, cruise, and descent. The model captures the variation in aerodynamic and propulsion requirements across

these phases and enables estimation of cumulative electrical energy consumption over the mission duration.

Total mission energy is computed as the time integral of electrical power:

$$E(t_{end}) = \frac{1}{3600} \int_0^{t_{end}} P_e(t) dt, \quad (1)$$

where $P_e(t)$ is the electrical power obtained from a two-dimensional (2D) lookup table parameterized by rotor speed and shaft torque and includes aggregated drivetrain losses [21]. The model accounts for variations in power demand resulting from changes in rotor speed, flight velocity, and trajectory profile.

3.3. Rotor Speed Scheduling Model

Rotor speed scheduling is implemented using phase-dependent setpoints corresponding to hover, ascent, cruise, and descent. The rotor speed profile is parameterized as:

$$X_\Omega = [\Omega_{hover} \ \Omega_{ascent} \ \Omega_{cruise} \ \Omega_{descent}]^T \quad (2)$$

The rotor speed commands are subject to actuator and dynamic constraints. A maximum slew rate constraint:

$$\left| \frac{d\Omega}{dt} \right| \leq 10 \text{RPM/s} \quad (3)$$

is imposed, and a first order filter:

$$G(s) = \frac{1}{\tau s + 1} \quad (4)$$

is applied to represent rotor dynamic response. This formulation enables evaluation of how phase-dependent rotor speed adjustments influence power consumption and mission energy.

3.4. Trajectory Optimization Formulation

The mission trajectory is parameterized using key flight variables, including phase-dependent true airspeeds (TAS) and cruise altitude:

$$X_V = [V_{ascent}, V_{cruise}, V_{descent}, h_{cruise}]^T, \quad (5)$$

Trajectory commands are constrained to ensure physically realistic transitions. Rate limits are imposed on velocity and altitude profiles:

$$\left| \frac{dV}{dt} \right| \leq 12 \text{ft/s}^2 \ (3.7 \text{ m/s}^2), \quad (6)$$

$$\left| \frac{dh}{dt} \right| \leq 5 \text{ft/s} \ (1.5 \text{ m/s}) \quad (7)$$

These constraints ensure smooth transitions between flight phases and maintain stable tracking performance within the nonlinear simulation environment.

3.5. Joint Optimization Framework

To capture the interaction between propulsion control and flight trajectory, a joint optimization framework is developed in which rotor speed scheduling and trajectory parameters are optimized simultaneously. The combined decision vector is defined as:

$$X = \begin{bmatrix} X_\Omega \\ X_V \end{bmatrix} \quad (8)$$

This formulation enables exploration of tradeoffs between aerodynamic loading and propulsion efficiency. For example, adjustments in climb rate or cruise altitude may allow lower rotor speed operation, resulting in reduced overall energy consumption.

3.6. Optimization Problem Formulation

The optimization problem is formulated as a constrained mission-based parametric search using the nonlinear rotorcraft simulation model.

The objective is to minimize total mission energy consumption:

$$\min_X J(X) = E(t_{end}), \quad (9)$$

subject to: rotor speed constraints, trajectory rate limits, flight dynamics feasibility, mission completion requirements.

The mission is divided into four phases with fixed durations: hover ($t_{hover} = 40$ s), ascent ($t_{ascent} = 30$ s), cruise ($t_{cruise} = 70$ s), and descent, completing the mission at $t = 200$ s.

Three optimization strategies are evaluated:

1. Rotor speed optimization with a fixed baseline trajectory, used to determine the optimal phase-dependent rotor speed schedule;
2. Trajectory optimization using the best optimal rotor speed schedule obtained from the previous stage and held constant for all cases;
3. Joint optimization of rotor speed and trajectory, performed to assess the integrated effect of both variables.

Candidate solutions are evaluated based on cumulative energy consumption, and only physically feasible trajectories with stable tracking behavior are accepted.

In contrast to the preliminary study presented in [1], the current work extends the analysis by incorporating a joint optimization framework and providing a systematic comparison of different optimization strategies. These extensions enable a

more comprehensive evaluation of the factors influencing the UAV rotorcraft mission energy efficiency.

3.7. Performance Metrics and Simulation Assumptions

Performance comparisons are conducted using both total cumulative energy consumption and flight phase energy distribution across hover, ascent, cruise, and descent segments. This allows assessment of how different optimization strategies redistribute energy demand throughout the mission profile, enabling identification of the dominant contributors to energy consumption.

To isolate the effects of rotor speed scheduling and trajectory optimization, all simulations are conducted under consistent environmental conditions. Standard atmospheric properties are assumed, with no wind or gust disturbances. This ensures that differences in energy consumption are attributable solely to variations in control inputs and trajectory parameters.

The UAV, propulsion and simulation parameters used in this study are summarized in Table 1, while the baseline mission profile is defined in Table 2. These parameters are consistent with those used in the preliminary study [1], with minor refinements introduced to support the extended analysis presented in this work.

Table 1. Rotorcraft and simulation parameters.

Parameter	Symbol	Value / Relation / Formula	Unit / Description
UAV mass	m	7.23	kg
Diameter of main rotor	D_m	1.78	m
Main rotor blade count	N_m	2	Number of blades (dimensionless)
Tail rotor diameter	D_t	0.158	m
Tail rotor blade count	N_t	2	Number of blades (dimensionless)
Baseline RPM	n_{base}	1000	RPM
Variable rotor speed range	n	850–1000	RPM
Rotor response dynamics	-	$G(s) = \frac{1}{\tau s + 1}$ (with rate limiting)	First order rotor speed dynamics
Electrical power map	P_e	2D LUT: $P_e(Q, n)$	W (electrical input power from powertrain; function of torque and speed)
Torque calculation	Q	From rotor model	N m (main rotor shaft torque)
Flight dynamics model	-	Nonlinear 6DOF rigid body model	Implemented in FLIGHTLAB
Trim methodology	-	Phase-wise steady state trim	Flight phases: hover, climb, cruise, descent
Energy model	E	$E(t_{end}) = \frac{1}{3600} \int_0^{t_{end}} P_e(t) dt$	Wh (time integration of electrical power P_e in W)
Atmospheric model	-	Standard atmosphere (ISA)	Assumes sea-level conditions, no environmental deviations
Wind / turbulence	-	Not considered	No wind or turbulence effects included

Table 2. Baseline mission profile parameters.

Phase	Start Time t_z (s)	Duration Δt_z (s)	Velocity V_z (ft/s [m/s])	Altitude H_z (ft [m])	Altitude Command
Hover	0	40	0	34	Constant
Ascent	40	30	60	80	Constant
Cruise	70	70	60	80	Climb rate, ± 2 ft/s
Descent	140	60	55	34	Target altitude 34 ft

4. Result and Discussion

4.1. Baseline Mission Energy Characteristics

The baseline mission represents the reference operating condition against which the effectiveness of the optimization strategies is evaluated. In this configuration, the rotorcraft operates at a constant rotor speed of 1000 RPM throughout the mission while following a fixed trajectory consisting of hover, ascent, cruise, and descent phases. The resulting total mission energy consumption for the baseline case is

36.06 Wh, which represents the highest energy demand among all evaluated configurations.

Fig. 2 illustrates the baseline mission trajectory, including altitude, forward airspeed, and vertical speed, with vertical markers indicating transitions between flight phases. The rotorcraft maintains a steady hover at 34 ft (10.4 m), followed by a gradual climb to the cruise altitude of approximately 80 ft (24.4 m). This altitude is sustained during the cruise segment before a controlled descent returns the UAV to the initial hover level.

The airspeed profile reflects the prescribed mission schedule shown in Table 2, increasing from zero in hover to the cruise value of approximately 60 ft/s (18.3 m/s), and then reducing slightly during descent at 55 ft/s (16.8 m/s). The vertical speed history further confirms the phase behavior, with positive climb rates during ascent, near zero values in steady segments, and negative rates during descent, indicating smooth and continuous transitions.

Overall, the time histories demonstrate consistent tracking of the commanded flight profile, with no abrupt transients or instability. This validates the baseline trajectory as a reliable reference case for subsequent power and energy analysis.




Fig. 2. Baseline mission flight profile showing altitude, forward airspeed, and vertical speed with indicated phase transitions between hover, ascent, cruise, and descent.

The baseline mission energy characteristics are illustrated using both time resolved and phase averaged representations. A constant rotor speed of 1000 RPM is used as a reference for evaluating optimization strategies.

Fig. 3 shows the time resolved main rotor (MR) power profile for the baseline mission. The hover phase exhibits the highest steady power demand approximately 591 W, dominated by induced power required to sustain lift. At the transition to ascent at 40s, a transient peak of approximately 635 W is observed due to the increased thrust required for climb. As forward speed develops, the power decreases progressively, reflecting improved aerodynamic efficiency associated with translational lift.

During cruise, the power stabilizes near 455 W, indicating steady aerodynamic conditions with reduced induced power and balanced drag contributions. At the onset of descent at 140 s, a sharp reduction in power occurs, followed by a brief transient dip, as gravitational effects reduce the required rotor thrust. The subsequent stabilization reflects controlled descent under lower power demand. These results highlight that hover and ascent phases dominate instantaneous power requirements.

The cumulative electrical energy consumption over the baseline mission is shown in Fig. 4, showing

a monotonic increase over time that reflects the continuous power demand of the propulsion system across all flight phases. The slope of the curve varies slightly between segments, indicating changes in power demand associated with different operating conditions.




Fig. 3. Baseline main rotor power profile over the mission timeline. Phase transitions are indicated, highlighting transient responses associated with changes in flight conditions.

During hover and ascent, the energy accumulation rate is relatively higher due to increased thrust requirements. A more gradual and nearly linear trend is observed during cruise, consistent with steady operating conditions. In the descent phase, the rate of increase remains moderate, reflecting reduced power

demand compared to earlier phases while still accounting for controlled flight.




Fig. 4. Cumulative electrical energy consumption for the baseline mission.

Fig. 5 provides a phase averaged view of power demand and cumulative energy consumption for the baseline case. The blue bars represent the average power in each phase, obtained by averaging the time-resolved MR power over the corresponding phase durations. Consistent with Fig. 2, hover exhibits the highest average power, approximately 611 W, followed by ascent, 527 W, while cruise and descent phases show lower values of approximately 456 W and 429 W, respectively. The reduction in average power across phases reflects the increasing contribution of translational aerodynamic effects and reduced thrust requirements.




Fig. 5. Flight phase average power and cumulative energy consumption for the baseline mission profile.

The red curve represents cumulative energy consumption per phase, which depends on both power level and phase duration. Although cruise does not exhibit the highest average power, it contributes the largest energy consumption 11.7 Wh because it occupies the longest duration in the mission. In contrast, ascent consumes less energy, approximately 5.9 Wh despite moderate power demand because of its shorter duration.

These results demonstrate that hover dominates instantaneous power demand, whereas cruise dominates total energy consumption due to mission duration. Consequently, reducing cruise power or shortening cruise duration offers the greatest potential for minimizing overall mission energy.

These baseline characteristics provide a reference for evaluating how rotor speed scheduling and

trajectory optimization modify both instantaneous power demand and cumulative energy consumption.

4.2. Effect of Rotor Speed Scheduling

This section evaluates the impact of rotor speed scheduling on mission energy consumption while maintaining the baseline trajectory. The rotor speed is adjusted across flight phases to improve aerodynamic efficiency and reduce power demand.

The impact of rotor speed scheduling on propulsion performance is first examined through the rotor speed time histories, as shown in Fig. 6. In the baseline configuration, the rotor speed is maintained constant at 1000 RPM throughout the mission. In contrast, the optimized schedule applies phase-dependent rotor speeds, with higher values during hover and ascent and reduced values during cruise and descent.

The optimized rotor speed decreases from approximately 925 RPM in hover to 870 RPM during cruise, reflecting a tradeoff between thrust requirements and aerodynamic efficiency. Lower rotor speeds in forward flight reduce profile power losses, while maintaining sufficient lift.

These results demonstrate that rotor speed scheduling enables adaptation of rotor operating conditions to changing aerodynamic requirements across flight phases, forming the basis for the observed reductions in power consumption and total mission energy.




Fig. 6. Rotor speed profiles for baseline and optimized cases showing reduced rotor speed across mission phases.

The resulting changes in power demand are illustrated in Fig. 7, which compares the main rotor power profiles for the baseline and optimized cases. The optimized cases exhibit a consistent reduction in power demand across all mission phases relative to the constant speed baseline.

The most significant reduction occurs during the cruise phase, where power decreases from approximately 460 W to about 350 W. This reduction is primarily attributed to lower rotor speed, which reduces profile power losses under forward flight conditions. Due to the extended duration of cruise, this sustained reduction contributes significantly to the overall energy savings.

During hover and ascent, power reductions are also observed, although the relative improvement is moderated by the dominance of induced power, which

limits the extent of achievable savings. Nevertheless, reduced rotor speed still leads to lower overall power demand while maintaining required thrust.

In the descent phase, power demand decreases further due to both reduced thrust requirements and lower rotor speed. The optimized cases maintain consistently lower power levels while preserving smooth transitions between flight phases.

These results demonstrate that rotor speed scheduling effectively reduces power consumption, particularly in phases where profile power dominates, thereby contributing to the observed reduction in total mission energy.




Fig. 7. Main rotor power comparison between the baseline and optimized cases showing lower power across mission phases.

The cumulative effect of these reductions is presented in Fig. 8, which shows the total energy consumption over time. All optimized cases exhibit a consistently lower rate of energy accumulation compared to the baseline, indicating sustained reductions in power demand across all flight phases.




Fig. 8. Cumulative energy consumption for baseline and optimized rotor speed schedules over the mission.

The separation between the baseline and optimized trajectories increases progressively with time, with the largest divergence occurring during the cruise phase, where reduced rotor speed lowers power demand over an extended period. The optimized cases converge to similar energy profiles, indicating that different rotor speed schedules can achieve comparable levels of energy efficiency.

The best performing case (V2) achieves a total energy reduction from approximately 36.06 Wh to 28.18 Wh, corresponding to an energy reduction of about 21.85 %, confirming the effectiveness of rotor speed scheduling in reducing mission energy consumption.

The energy consumption distribution across individual flight phases is summarized in Fig. 9, which

compares the energy required for each segment. Energy reductions are observed across all flight phases relative to the baseline. In hover, energy consumption decreases from approximately 9.0 Wh to 7.0 Wh, while in ascent it reduces from 5.9 Wh to 4.6 Wh. During cruise, the energy decreases from 11.7 Wh to 9.2 Wh, representing the largest absolute reduction due to the longer duration of this phase. A similar reduction is observed during descent, from 9.5 Wh to 7.4 Wh.

The results discussed in this section establish a clear relationship between rotor speed scheduling, reduced power demand, and improved energy efficiency. Lower rotor speeds reduce aerodynamic losses, particularly in forward flight, leading to sustained reductions in cumulative energy consumption across the mission.




Fig. 9. Energy consumption across mission phases for the baseline and selected rotor speed schedules showing reduced total energy for the optimal case.

To provide a comprehensive overview of the evaluated rotor speed schedules, the results for all tested configurations are summarized in Table 3. A total of six candidate schedules were assessed in addition to the baseline case. The results show that all optimized configurations achieve significant reductions in total mission energy, with reductions exceeding 20 % across all cases.

4.3. Trajectory Optimization

In this section, trajectory optimization is performed within a sequential optimization framework using a fixed rotor speed schedule obtained from the preceding rotor speed optimization stage. The rotor speed profile (925, 870, 870, 850) is held constant for all optimized cases, allowing the influence of trajectory parameters on energy consumption to be isolated. Specifically, the ascent speed, cruise speed, descent speed, and cruise altitude are varied, while maintaining a consistent rotor operating condition across all simulations. By modifying these key parameters, the vehicle can

operate under more favorable aerodynamic conditions, thereby reducing overall energy demand.

The impact of trajectory design on mission energy consumption is evaluated by comparing the baseline case with two representative optimized trajectories, as

illustrated in Figs. 10–13. The selected cases include the baseline profile (R0), the best performing optimized trajectory (R1), and a higher altitude alternative (R2), enabling assessment of both speed and altitude effects.

Table 3. Summary of rotor speed schedules and phase energy consumption across the mission.

Case	RPM Schedule	Total Energy (Wh)	Energy Reduction (%)	Phase Energy (Wh)			
				Hover	Ascent	Cruise	Descent
Baseline	(1000, 1000, 1000, 1000)	36.06	0.00	8.99	5.85	11.75	9.47
V0	(900, 850, 855, 850)	28.27	21.58	7.02	4.60	9.24	7.42
V1	(880, 830, 840, 835)	28.55	20.81	7.05	4.65	9.38	7.47
V2 (Best)	(925, 870, 870, 850)	28.18	21.85	7.00	4.57	9.20	7.41
V3	(925, 875, 870, 860)	28.20	21.78	7.00	4.57	9.21	7.43
V4	(900, 850, 845, 840)	28.37	21.33	7.02	4.60	9.30	7.44
V5	(860, 830, 855, 850)	28.49	20.97	7.09	4.66	9.32	7.42

The altitude profiles shown in Fig. 10 indicate that all trajectories follow similar mission structures with distinct phase transitions. The baseline and optimized case R1 maintain a cruise altitude of approximately 80 ft (24.4 m), while R2 operates at a higher cruise altitude of 90 ft (27.4 m). The climb and descent segments remain smooth across all cases, ensuring comparable mission feasibility. However, R1 exhibits a slightly earlier and more gradual descent initiation, which contributes to improved energy efficiency.




Fig. 10. Altitude profiles for baseline and optimized trajectory cases.

The corresponding airspeed profiles in Fig. 11 reveal the primary source of variation between trajectories. The baseline maintains a cruise speed of approximately 60 ft/s (18.3 m/s), whereas R1 adopts a reduced speed profile 50, 45, and 40 ft/s (15.2, 13.7, and 12.2 m/s), and R2 increases cruise speed to approximately 65 ft/s (19.8 m/s). These differences directly influence aerodynamic power requirements, particularly during the cruise phase where forward flight dominates energy consumption.

The effect of these trajectory modifications is clearly reflected in the main rotor power profiles presented in Fig. 12. The optimized trajectory R1 achieves a consistent reduction in power demand across all phases, with the most significant improvement observed during cruise. In this phase, the power requirement decreases substantially compared to the baseline, indicating that reduced forward speed

effectively lowers profile and parasitic power components. In contrast, the higher altitude case R2 exhibits moderate power reduction relative to the baseline but remains less efficient than R1 due to increased aerodynamic loading associated with higher cruise speed.




Fig. 11. Airspeed profiles showing variations in optimized trajectories.

The cumulative impact of these power reductions is shown in Fig. 13, where the total energy consumption over the mission is compared. The baseline case reaches approximately 36.06 Wh, while the optimized trajectory R1 reduces the total energy to 27.05 Wh, corresponding to an energy reduction of 24.99 %. The higher altitude case R2 achieves a smaller reduction of 19.74 %, confirming that altitude increase alone does not yield optimal performance without coordinated speed adjustment.




Fig. 12. Main rotor power comparison for baseline and optimized trajectories illustrating reduced power demand under optimized trajectory conditions.




Fig. 13. Cumulative energy consumption showing 24.99 % reduction achieved through trajectory optimization.

These results demonstrate that trajectory optimization plays a critical role in reducing mission energy consumption. Specifically, reducing forward airspeed during energy intensive phases, particularly cruise, leads to substantial improvements in efficiency. While altitude variations influence performance, the results indicate that speed optimization is the dominant factor. The best performing trajectory (R1) achieves the lowest energy consumption by balancing reduced aerodynamic power demand with feasible flight dynamics, thereby establishing an effective strategy for mission-based energy optimization for electric UAV rotorcraft.

The results of the trajectory optimization study are summarized in Table 4. The variations in trajectory parameters lead to significant differences in total mission energy, despite identical rotor speed profiles.

The baseline case, defined by a speed profile of 60, 60, and 55 ft/s (18.3, 18.3, and 16.8 m/s) and a cruise altitude of 80 ft (24.4 m), with a constant rotor speed of 1000 RPM, results in a total energy consumption of 36.06 Wh. In contrast, all optimized cases are evaluated using the prescribed rotor speed schedule, ensuring that only trajectory parameters are varied.

The optimized trajectories achieve substantial reductions, with energy reductions ranging from 19.74 % to 24.99 %. The best performing case, R1, reduces total energy to 27.05 Wh, corresponding to a 24.99 % improvement, primarily due to lower forward speeds during ascent, cruise, and descent. In contrast, higher speed or higher altitude configurations, such as R2 and R3, provide smaller benefits despite increased maximum true airspeed (TAS), indicating that higher aerodynamic loads offset potential efficiency gains. The reduced altitude case R4 also performs well, achieving a 23.06 % reduction, indicating that both speed and altitude influence energy consumption, though speed has a more dominant effect. Across all cases, maximum climb and descent rates remain within a narrow range, confirming that performance constraints are maintained. These results indicate that trajectory optimization, when performed under the fixed rotor speed schedule, provides substantial energy savings and represents a dominant contributor to overall mission efficiency within the proposed sequential optimization framework.

Table 4. Summary of trajectory optimization results obtained using a fixed rotor speed schedule derived from prior rotor speed optimization. The rotor speed profile is held constant for all optimized cases.

Case	Fixed Rotor Speed Schedule (RPM)	Trajectory Parameters (V _{ascent} , V _{cruise} , V _{descent} (ft/s [m/s]); h _{cruise} (ft [m]))	Total Energy (Wh)	Energy Reduction (%)	Max TAS (ft/s [m/s])	Max Climb (ft/s [m/s])	Max Descent (ft/s [m/s])
R0 (Baseline)	(1000, 1000, 1000, 1000)	(60, 60, 55 ft/s; 80 ft)	36.06	0.00	59.4	2.94	-1.85
R1 (Best)	(925, 870, 870, 850)	(50, 45, 40 ft/s; 80 ft)	27.05	24.99	54.9	2.18	-1.88
R2	(925, 870, 870, 850)	(60, 65, 60 ft/s; 90 ft)	28.94	19.74	64.1	2.18	-1.90
R3	(925, 870, 870, 850)	(60, 60, 55 ft/s; 90 ft)	28.18	21.85	59.4	2.18	-1.89
R4	(925, 870, 870, 850)	(60, 55, 50 ft/s; 70 ft)	27.74	23.06	59.0	2.18	-1.81
R5	(925, 870, 870, 850)	(60, 60, 60 ft/s; 80 ft)	28.52	20.89	59.7	2.18	-1.94

4.4. Joint Optimization Analysis

This section evaluates the simultaneous optimization of rotor speed scheduling and flight trajectory to assess their combined impact on mission energy consumption. While the previous sections examined these strategies independently, their interaction can produce additional performance gains that are not achievable through isolated optimization. By jointly adjusting rotor speed and trajectory parameters, the system can better balance aerodynamic and rotor power requirements across all flight phases. The following analysis compares the joint approach with the baseline and individually optimized cases, with emphasis on total energy consumption, power distribution, and operational efficiency.

A parametric search was conducted by integrating optimized trajectory parameters with varying rotor speed schedules, as summarized in Table 5. The baseline case employs a constant rotor speed, whereas the combined cases (C1–C6) implement reduced rotor speeds alongside optimized flight path parameters. All combined configurations demonstrate significant reductions in total energy compared to the baseline.

Among the evaluated cases, C1 achieves the lowest total energy consumption of 27.04 Wh, corresponding to a 25.01 % reduction relative to the baseline. Notably, cases C2, C3 and C4 yield nearly identical energy values despite slight variations in rotor speed schedules, indicating the presence of a relatively flat optimum region. This behavior suggests that the combined solution is robust to moderate changes in rotor speed when the trajectory parameters are maintained.


Table 5. Joint optimization results for rotor speed and trajectory parameters.

Case	RPM Schedule	Trajectory Parameters (V _{ascent} , V _{cruise} , V _{descent} (ft/s [m/s]); h _{cruise} (ft [m]))	Total Energy (Wh)	Energy Reduction (%)
Baseline	(1000, 1000, 1000, 1000)	(60, 60, 55 ft/s; 80 ft)	36.06	0.00
C1 (Best)	(935, 880, 870, 850)	(50, 45, 40 ft/s; 80 ft)	27.04	25.01
C2	(925, 870, 870, 850)	(50, 45, 40 ft/s; 80 ft)	27.05	24.99
C3	(915, 860, 860, 845)	(50, 45, 40 ft/s; 80 ft)	27.05	24.97
C4	(935, 880, 880, 855)	(50, 45, 40 ft/s; 80 ft)	27.05	24.99
C5	(925, 870, 870, 850)	(50, 50, 40 ft/s; 80 ft)	27.20	24.57
C6	(925, 870, 870, 850)	(50, 45, 45 ft/s; 80 ft)	27.15	24.70

In contrast, Cases C5 and C6 introduce small deviations from the optimal trajectory by modifying the cruise and descent speeds, respectively. In C5, the cruise speed is increased from 45 ft/s (13.7 m/s) to 50 ft/s (15.2 m/s), leading to higher aerodynamic power during forward flight and a corresponding increase in energy consumption. In C6, the descent speed is increased from 40 ft/s (12.2 m/s) to 45 ft/s (13.7 m/s), reducing the efficiency of the low power descent phase. These deviations result in slightly higher total energy values compared to the best case, highlighting the sensitivity of the solution to forward speed selection, particularly during cruise.

The underlying control mechanism of the combined solution is illustrated in Fig. 14 which


compares the baseline constant rotor speed with the optimized schedule. The baseline maintains a fixed value of 1000 RPM throughout the mission, whereas the optimized schedule adapts to the trajectory. Reduced rotor speeds are maintained during cruise and descent, where aerodynamic loading is lower, while slightly higher values are preserved during hover and ascent to satisfy thrust requirements. This variation reflects the changing balance between induced and profile power, where lower rotor speeds are advantageous in forward flight conditions. The smooth transitions between phases also indicate that the optimized schedule remains physically realizable and avoids abrupt control inputs.

**Fig. 14.** Rotor speed schedule for baseline and joint optimization across flight phases.

The impact of this optimized control strategy on power demand is shown in Fig. 15. The combined case reduces power demand across all flight phases, with the most pronounced reductions occurring during cruise and transition periods. In these phases, the simultaneous reduction of rotor speed and forward airspeed lowers both profile and parasitic power components, leading to a more efficient operating condition. Compared to the individual optimization cases, the joint strategy achieves a more uniform

reduction in power, indicating improved overall system efficiency.

The improvement achieved through the joint optimization is illustrated in Fig. 16, where cumulative energy consumption is compared across the baseline, individual optimization cases, and the joint approach. The joint optimization consistently exhibits the lowest energy accumulation throughout the mission, indicating that the benefits of rotor speed reduction and trajectory shaping are complementary rather than redundant.

**Fig. 15.** Main rotor power comparison for baseline, rotor speed optimization, trajectory optimization, and joint optimization.




Fig. 16. Cumulative energy consumption for baseline, rotor speed optimization, trajectory optimization, and joint optimization.

The quantitative comparison of the different optimization strategies is summarized in Table 6. The baseline case requires a total energy of 36.06 Wh. Rotor speed optimization reduces this to 28.18 Wh, corresponding to a 21.85 % reduction. Trajectory

optimization provides a greater improvement, lowering total energy to 27.05 Wh (24.99 %), which represents an additional 4.01 % improvement relative to rotor speed optimization.

Table 6. Total mission energy consumption and relative performance comparison of rotor speed optimization, trajectory optimization, and joint optimization.

Case	Total Energy (Wh)	Energy Saving (%)	Improvement vs Rotor Speed Optimization (%)	Improvement vs Trajectory Optimization (%)
Baseline	36.006	0.00	-	-
Rotor speed optimization	28.18	21.85	0.00	-4.18
Trajectory Optimization	27.05	24.99	4.01	0.00
Joint Optimization	27.04	25.01	4.05	0.04

The joint optimization achieves the lowest total energy consumption of 27.04 Wh, corresponding to a 25.01 % reduction from the baseline. This represents a 4.05 % improvement relative to rotor speed optimization and a marginal additional improvement of 0.04 % over trajectory optimization. Given the small magnitude of this difference, the additional gain may fall within numerical tolerance or model uncertainty and should be interpreted with caution.

These results indicate that trajectory optimization provides the dominant contribution to energy reduction, while rotor speed scheduling offers incremental benefits when combined with trajectory optimization. However, further analysis, such as convergence or sensitivity studies, influence of environmental and operational uncertainties would be required to establish the robustness of this additional gain. The improvements are primarily driven by reduced power demand during cruise, with additional contributions from optimized ascent and descent profiles.

All simulations are performed under standard atmospheric conditions without explicitly considering wind, turbulence, or parameter uncertainty. These assumptions were intentionally adopted to provide a controlled environment for isolating the effects of rotor speed scheduling and trajectory optimization on mission energy consumption.

Variations in environmental conditions, vehicle mass, and battery state are expected to influence the absolute energy values reported. In particular, increased mass would primarily affect induced power during hover and ascent, while wind conditions would alter optimal trajectory parameters, especially during cruise. However, the core findings of this study are based on the interaction between rotor speed and aerodynamic loading, which governs the relative trends observed across optimization strategies. These underlying physical relationships are expected to remain consistent under moderate variations in operating conditions.

Therefore, while absolute energy savings may vary, the relative performance and benefits of the joint optimization framework are expected to remain qualitatively robust. Future work will extend the present analysis by incorporating environmental disturbances, parameter uncertainty, and adaptive mission planning under more realistic operating conditions.

5. Conclusions

This study presented an integrated framework for minimizing mission energy consumption in an electric rotorcraft through rotor speed scheduling and trajectory optimization. A physics-based simulation

model was used to evaluate the influence of rotor speed and trajectory parameters on power demand and cumulative energy consumption over a complete mission profile.

The results show that trajectory optimization provides the dominant contribution to energy reduction, while rotor speed scheduling offers complementary improvements within the sequential optimization framework. The rotor speed scheduling achieves a reduction of 21.85 %, while trajectory optimization yields a larger reduction of 24.99 % relative to the baseline. When both strategies are applied simultaneously, the combined optimization achieves the lowest total energy consumption of 27.04 Wh, corresponding to an overall reduction of 25.01 %.

Analysis of rotor speed, main rotor power, and cumulative energy profiles shows that the primary mechanism for energy reduction is the decrease in power demand during cruise, where reduced forward speed and lower rotor speed minimize aerodynamic losses. The joint optimization further improves efficiency by coordinating rotor speed with trajectory changes, resulting in consistent power reductions across all flight phases.

Although the additional improvement of combined optimization over trajectory optimization alone is modest, approximately 0.04 %, this difference may fall within numerical tolerance or model uncertainty and should therefore be interpreted with caution. The results nevertheless indicate consistent behavior between rotor speed scheduling and trajectory optimization, suggesting that the joint approach operates within a region close to the minimum achievable energy, with diminishing returns beyond trajectory optimization. The observed behavior also indicates a relatively flat optimum region, implying that comparable energy performance can be achieved under moderate variations in control parameters.

Overall, the findings highlight the importance of jointly considering trajectory and rotor speed in the design of energy-efficient rotorcraft operations. The proposed approach provides a systematic methodology for evaluating mission-based energy performance and can be extended to more complex scenarios, including variable environmental conditions, additional constraints, and real-time optimization strategies.



Acknowledgements

The authors express their gratitude for the computational facilities and simulation tools that made this research possible.

References

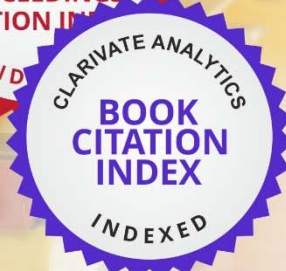

- [1]. E. Chijioke, M. Žugaj, Comparative and joint optimization of rotor speed and trajectory for energy efficient electric Unmanned Aerial Vehicle (UAV) helicopters, in *Proceedings of the 2nd International Conference on Drones and Unmanned Systems (DAUS'26)*, 2026, pp. 191-197.
- [2]. M. Shamiyeh, J. Bijewitz, M. Hornung, A review of recent personal air vehicle concepts, in *Proceedings of the Aerospace Europe 6th CEAS Conference*, 2017, pp. 1-18.
- [3]. M. D. Moore, B. Fredericks, Misconceptions of electric propulsion aircraft and their emergent aviation markets, in *Proceedings of the 52nd Aerospace Sciences Meeting*, 2014, 0535.
- [4]. B. J. Brelje, J. R. R. A. Martins, Electric, hybrid, and turboelectric fixed-wing aircraft: A review of concepts, models, and design approaches, *Progress in Aerospace Sciences*, Vol. 104, 2019, pp. 1-19.
- [5]. J. G. Leishman, Principles of Helicopter Aerodynamics (2nd Ed.), *Cambridge University Press*, New York, 2006.
- [6]. W. Johnson, Rotorcraft Aeromechanics, *Cambridge University Press*, New York, 2013.
- [7]. W. Garre, T. Pflumm, M. Hajek, Enhanced efficiency and flight envelope by variable main rotor speed for different helicopter configurations, in *Proceedings of the 42nd European Rotorcraft Forum*, 2016, pp. 722-735.
- [8]. M. Mistry, F. Gandhi, Helicopter performance improvement with variable rotor radius and RPM, *Journal of the American Helicopter Society*, Vol. 59, Issue 4, 2014, 042010.
- [9]. D. Han, V. Pstrikakis, G. N. Barakos, Helicopter performance improvement by variable rotor speed and variable blade twist, *Aerospace Science and Technology*, Vol. 54, 2016, pp. 164-173.
- [10]. H. Rienecker, V. Hildebrand, H. Pflifer, Energy optimal 3D flight path planning for unmanned aerial vehicle in urban environments, *CEAS Aeronautical Journal*, Vol. 14, 2023, pp. 621-636.
- [11]. L. Deori, S. Garatti, M. Prandini, A model predictive control approach to aircraft motion control, in *Proceedings of the American Control Conference (ACC'15)*, 2015, pp. 2299-2304.
- [12]. G. Su, X. Cheng, X. Liu, J. Jiang, Flight path optimization with optimal control method, *arXiv*, 2024, arXiv:2405.08306.
- [13]. E. C. Suicmez, A. T. Kutay, Optimal path tracking control of a quadrotor UAV, in *Proceedings of the International Conference on Unmanned Aircraft Systems (ICUAS'14)*, 2014, pp. 115-125.
- [14]. D. C. Gandolfo, L. R. Salinas, A. S. Brandão, J. M. Toibero, Path following for unmanned helicopter: An approach on energy autonomy improvement, *Information Technology and Control*, Vol. 45, Issue 1, 2016, pp. 86-98.
- [15]. A. M. Mendoza, R. M. Botez, Commercial aircraft trajectory optimization to reduce flight costs and pollution: Metaheuristic algorithms, in *Advances in Visualization and Optimization Techniques for Multidisciplinary Research* (D. Vucinic, F. R. Leta, S. Janardhanan, Eds.), *Springer*, Singapore, 2020, pp. 33-62.
- [16]. H. Kasmi, S. Laporte, R. M. Botez, Holistic approach for aircraft trajectory optimization using optimal control, *Journal of Aircraft*, Vol. 60, Issue 4, 2023, pp. 1302-1313.
- [17]. A. Visintini, T. D. P. Perera, A. V. Savkin, A. Zanella, 3-D trajectory optimization for fixed-wing UAV-enabled wireless network, *IEEE Access*, Vol. 9, 2021, pp. 35045-35056.

- [18]. M. A. Sadi, A. Jamali, N. M. N. Al-Falahi, M. F. Hossain, et al., Cascade model predictive control for enhancing UAV quadcopter stability and energy efficiency in wind turbulent mangrove forest environment, *e-Prime – Advances in Electrical Engineering, Electronics and Energy*, Vol. 10, 2024, 100836.
- [19]. I. Nagy, E. Laufer, Energy-optimized 3D path planning for unmanned aerial vehicles, *Applied Sciences*, Vol. 14, Issue 16, 2024, 6988.
- [20]. F. Yacef, N. Rizoug, L. Dala, C. Machado, Energy-efficiency path planning for quadrotor UAV under wind conditions, in *Proceedings of the 7th International Conference on Control, Decision and Information Technologies (CoDIT'20)*, 2020, pp. 1133-1138.
- [21]. M. Żugaj, M. Edawdi, G. Iwański, S. Topczewski, et al., An unmanned helicopter energy consumption analysis, *Energies*, Vol. 16, Issue 4, 2023, 2067.



10 Top Reasons to Publish Open Access Books with IFSA Publishing

- Indexed in Book Citation Index (Web of Science)
- Copyrights belong to Authors (CC-BY)
- The maximum number of pages is not limited
- Very reasonable publication fees
- High visibility
- All book types accepted
- Available in different formats: electronic and print
- Freely available online
- High quality standards
- Authors benefit from IFSA Membership



A Transport-Layer-Oriented Servo Actuation Framework for UAV Systems Using USB-CAN Communication

^{1,*} Jinyoung LEE and ² Sungwook CHO

¹ Cheongju University, Department of Mechanical and Aeronautical Systems Engineering,
298 Daeseong-ro, Cheongwon-gu, Cheongju-si, Chungbuk, South Korea

² Cheongju University, Department of Aeronautical and Mechanical Engineering,
298 Daeseong-ro, Cheongwon-gu, Cheongju-si, Chungbuk, South Korea

¹ Tel.: + 82 43 229 7948,

* E-mail:swcho84@cju.ac.kr

Received: 30 Jan. 2026 /Revised: 13 April 2026 /Accepted: 21 April 2026 /Published: 28 April 2026

Abstract: This study introduces a lightweight software framework for controlling servo-based actuation systems in unmanned aerial platforms through a universal serial bus to controller area network communication interface. The proposed approach emphasizes reliable command transmission, structured feedback acquisition, and deterministic communication behavior under general-purpose operating environments. Instead of focusing on closed-loop control design, the framework addresses challenges related to packet framing, timing consistency, and data integrity during continuous communication between host systems and actuator units. A custom byte-level framing structure is employed to ensure clear boundary detection, error validation, and stable parsing of actuator state information such as position, velocity, and electrical current. The system adopts a dual-layer scheduling mechanism to separate command execution from feedback collection, improving temporal consistency and reducing communication interference. Experimental validation demonstrates that the proposed framework achieves stable actuation performance with minimal deviation between commanded and measured states. The results confirm that the system provides a practical and extensible foundation for reliable actuator communication in aerial robotic platforms.

Keywords: Servo motor control, CAN communication protocol, Real-time control loop, Concurrent process scheduling, Modular software framework.

1. Introduction

Servo-driven actuation systems constitute fundamental components in unmanned aerial platforms, enabling precise mechanical responses in applications such as payload stabilization, camera orientation, and aerodynamic control surface adjustment. The performance of these actuation units directly affects the overall stability, sensing accuracy, and mission reliability of aerial systems, particularly under dynamic operating conditions involving rapid motion or environmental disturbances.

In conventional implementations, servo actuation is typically realized through pulse-width modulation interfaces or higher-level middleware frameworks that

abstract device communication and control processes. While such approaches offer simplicity and broad compatibility, they inherently limit direct access to low-level communication behavior, including timing consistency, data integrity, and synchronization between command transmission and feedback acquisition. As a result, system-level visibility into communication-induced effects such as latency variation, frame misalignment, or feedback inconsistency remains restricted. In many practical configurations, servo actuators internally execute closed-loop control strategies, whereas external host systems primarily provide reference inputs and interpret feedback data. Under these conditions, the reliability of the communication layer emerges as a

critical factor influencing overall system performance. Variability in transmission timing, buffering behavior, and data reconstruction can directly impact the consistency of actuator responses, even when the underlying control mechanisms remain stable [1, 2].

Motivated by these limitations, this work focuses on the design of a lightweight communication framework that prioritizes deterministic data exchange and structured timing control in servo-based actuation systems. Rather than proposing a new control algorithm, the proposed approach emphasizes the robustness of command delivery and the integrity of telemetry acquisition under non-real-time operating environments. A USB–CAN interface is employed to facilitate communication between the host system and actuator devices, while a custom byte-structured framing mechanism is introduced to enable explicit boundary detection and reliable parsing of continuous data streams [3].

To further enhance temporal consistency, the system adopts a dual-loop scheduling strategy that separates command transmission from telemetry polling processes. This design reduces communication contention and enables stable interaction between control inputs and feedback data, even in the presence of operating system-induced timing variability. The resulting framework is designed to provide transparency in communication behavior, allowing direct observation and analysis of transport-layer characteristics that are often abstracted in conventional middleware-based systems.

The contributions of this work can be summarized as follows. First, a communication-oriented servo actuation framework is developed to ensure consistent and verifiable data exchange between host systems and actuator units. Second, a custom framing scheme is introduced to improve frame synchronization and data validation in continuous serial communication environments. Third, a dual-loop scheduling mechanism is implemented to maintain stable timing behavior under non-deterministic system conditions. Finally, the effectiveness of the proposed approach is demonstrated through experimental evaluation, highlighting its suitability for integration into unmanned aerial systems requiring reliable and transparent actuator communication.

This work extends our previous conference paper presented at DAUS 2026. While the earlier study provided a preliminary implementation of a USB–CAN-based servo communication system, the present work significantly expands the framework in several aspects. First, a transport-layer-oriented architecture is newly formalized, explicitly addressing byte-level framing, data validation, and communication robustness. Second, a dual-loop scheduling mechanism is introduced to improve temporal consistency under non-real-time operating conditions. Third, a detailed system architecture is presented, including modular layer separation and communication pipeline design. Finally, the experimental evaluation is extended to include quantitative analysis of latency, timing variability, and

communication stability. These extensions distinguish the present work from the initial conference version and provide a more comprehensive and systematic framework for reliable actuator communication in UAV systems, moving beyond a preliminary implementation toward a rigorously structured communication architecture. These improvements enable a more rigorous evaluation of communication-layer behavior and provide stronger experimental evidence compared to the initial conference version [4].

2. Related Works

2.1. Conventional Servo Actuation Interfaces

Servo actuation in unmanned aerial systems has traditionally been implemented using pulse-width-modulation interfaces due to their simplicity and widespread hardware compatibility. However, such approaches provide limited support for structured telemetry, fault monitoring, and communication-level validation. In many cases, actuator state information such as position, velocity, and current is not directly accessible, which restricts observability and reduces suitability for applications requiring reliable state feedback and diagnostic capability.

2.2. Middleware-Based Robotic Frameworks

Modern robotic systems increasingly adopt middleware architectures such as Robot Operating System 2 to enable modular integration and distributed processing. Despite these advantages, recent studies have shown that communication latency and timing consistency in such frameworks are strongly influenced by the underlying operating system and execution model. In particular, non-real-time environments may introduce variability in message delivery, leading to jitter and reduced temporal predictability. These characteristics can significantly affect actuator communication, where consistent timing and reliable feedback acquisition are critical for stable system behavior.

2.3. CAN-Based Distributed Communication

Controller area network-based communication frameworks, including Cyphal, provide robust and scalable solutions for distributed embedded systems. These approaches emphasize interoperability, fault tolerance, and structured message exchange across multiple nodes. However, they typically abstract transport-layer details such as byte-stream handling, frame boundary reconstruction, and host-side scheduling behavior. While this abstraction is beneficial for large-scale integration, it limits direct visibility into communication timing and data integrity at the host-to-actuator interface [5-7].

2.4. Research Motivation

Existing approaches can be broadly categorized into pulse-width-modulation interfaces with limited observability, middleware-driven frameworks that offer strong integration but abstract timing behavior, and controller area network-based ecosystems designed for distributed communication. However, a critical gap remains in the design of lightweight actuator communication frameworks that explicitly expose transport-layer behavior – such as frame synchronization, payload validation, and timing control – while maintaining reliable and consistent

timing under non-real-time conditions. To clarify these differences, the key characteristics of representative actuation communication frameworks are summarized in Table 1, focusing on communication interface, data integrity, temporal control, and system observability. As shown in Table 1, conventional approaches either provide limited visibility into actuator behavior or abstract communication mechanisms, whereas the proposed framework explicitly exposes timing structure and frame-level validation, enabling improved communication transparency and more predictable actuator behavior in non-real-time conditions [7].

Table 1. Comparative Analysis of Communication Characteristics in UAV Actuation Systems.

Feature	PWM Interface	CAN-Based Middleware	ROS2-Based Control	Proposed Framework
Communication Interface	Pulse Signal	CAN Bus	Middleware Messaging	USB–CAN Communication
Telemetry Availability	Limited	Structured	Structured	Structured (multi-parameter)
Data Integrity Mechanism	None	Protocol-Level Validation	Middleware-Level Handling	Explicit Frame Validation
Temporal Control (Host-Side)	Fixed	Abstracted	Execution-Dependent	Explicit Scheduling
Update Rate Control	Fixed	Configurable	Configurable	User-Defined
System Observability	Low	Medium	Medium	High
Integration Complexity	Low	Medium	High	Medium

3. System Architecture

3.1. Overall Framework Design

The proposed system is organized as a layered communication and control framework that separates command generation, data transmission, and actuator interaction to ensure stable and predictable operation. Instead of directly coupling user-level commands with hardware communication, the architecture introduces a modular structure in which each functional layer operates independently while maintaining synchronized data exchange. At the top level, the control layer is responsible for generating actuator commands and managing periodic execution. This layer schedules reference updates at a fixed interval and stores outgoing messages in a concurrent buffer structure, allowing safe and consistent interaction with lower-level processes. By isolating command generation from communication handling, the system reduces timing interference caused by user interaction or high-level logic.

The communication layer operates as an intermediary that translates scheduled commands into a structured byte stream and manages both transmission and reception processes. This layer performs frame construction, parsing, and logging, ensuring that all data exchanged with the actuator system is consistently formatted and traceable. It also maintains continuous monitoring of incoming data,

enabling immediate detection of malformed or incomplete frames.

The actuator layer consists of a servo motor unit equipped with an internal driver and encoder-based feedback mechanism. This unit executes position commands and continuously provides telemetry data, including position, velocity, and current measurements. The separation of responsibilities between host-side communication and actuator-side control enables the system to focus on reliable data transport rather than control algorithm implementation.

3.2. Layered Functional Structure

To improve modularity and timing robustness, the system is divided into four functional layers: application, scheduling, communication, and data interpretation. Each layer is designed to operate independently while exchanging information through well-defined interfaces.

The application layer provides user interaction and visualization capabilities, allowing real-time monitoring of actuator states and manual or scripted command input. This layer does not directly interact with hardware, which prevents user-side operations from affecting communication timing.

The scheduling layer manages periodic task execution and prioritizes command transmission over

background processes. It enforces a structured update cycle and coordinates access to shared resources, ensuring that command delivery remains consistent even when telemetry polling is active.

The communication layer handles message encoding and decoding over the serial interface. It constructs outgoing frames, transmits them through the interface module, and reconstructs incoming frames from continuous byte streams. This layer also maintains transmission and reception logs, enabling post-analysis of communication behavior.

Finally, the data interpretation layer converts raw payload data into engineering units and distributes processed information to external clients. This separation allows telemetry data to be analyzed and visualized without interfering with the communication pipeline.

3.3. Command Scheduling and Arbitration Strategy

To maintain stable communication behavior, the system employs a structured scheduling and arbitration mechanism that separates time-critical command transmission from background telemetry acquisition. Rather than allowing simultaneous access to the communication channel, the scheduler enforces a priority-based policy in which actuator commands are always transmitted before non-essential requests.

During active command updates, telemetry polling is temporarily suspended to prevent contention on the communication channel. This approach ensures that control commands are delivered within the intended time window, reducing the risk of delayed or skipped updates. Once the command transmission interval is satisfied, telemetry requests are resumed, enabling continuous monitoring of actuator state without compromising control stability.

In addition, a periodic refresh mechanism is implemented to maintain actuator engagement during idle periods. When no new command is issued within a predefined interval, the system retransmits the most recent valid command to prevent unintended actuator relaxation or drift. This strategy enhances reliability by ensuring that the actuator consistently maintains its last commanded state even in the absence of continuous input. The overall structure of the proposed framework is illustrated in Fig. 1.

4. AA-55 Frame Structure

4.1. Frame Design Overview

To ensure reliable communication over a continuous serial data stream, the proposed framework employs a lightweight byte-oriented framing structure referred to as the AA-55 format. This design enables explicit identification of frame boundaries and

consistent reconstruction of messages, even under non-deterministic buffering conditions.

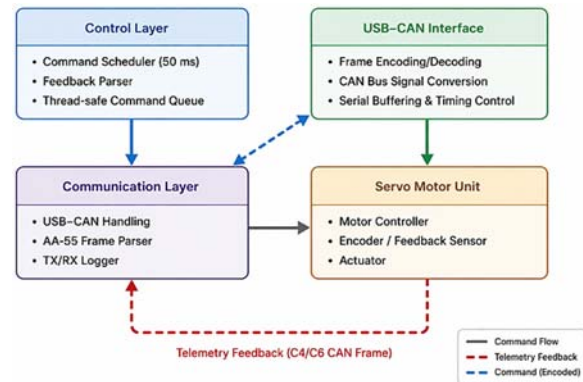


Fig. 1. Proposed Servo Communication Framework.

Each transmitted message is encapsulated within a structured frame composed of a start marker, control fields, payload data, and an end marker. The frame structure follows a fixed format: [Start] [Command] [ID] [Length] [Payload] [Checksum] [End]. The use of fixed boundary indicators allows the receiver to detect frame transitions without relying on external synchronization mechanisms. This approach is particularly effective in environments where partial reads or byte-level misalignment may occur during serial communication.

In addition to boundary detection, the frame structure incorporates a simple validation mechanism to ensure payload consistency. By verifying both the expected payload length and the integrity field, the receiver can identify malformed or incomplete frames before higher-level processing. This contributes to stable communication behavior by preventing corrupted data from propagating through the system.

4.2. Frame Field Composition

The AA-55 frame is organized into a sequence of fields that define both structural and semantic information. Each frame begins with a fixed start marker and ends with a corresponding termination marker, forming a clear boundary for message reconstruction. Between these markers, control and data fields specify the purpose and content of the message.

The command identifier indicates the type of operation, such as motion control or telemetry request, while the identifier field specifies the target actuator or parameter group. The payload length field defines the size of the data segment, enabling consistency checks during parsing. The data field contains actuator-related information, including position, velocity, current, and system status values. Finally, an integrity field provides a lightweight validation

mechanism that allows early detection of transmission errors at the host level.

This structured representation enables a unified handling of both command and feedback messages, simplifying the communication pipeline while maintaining flexibility for different actuator operations.

4.3. Frame Reconstruction Process

Incoming data from the serial interface are processed as a continuous byte stream without predefined message boundaries. To reconstruct valid frames, the receiver employs a state-based parsing mechanism that continuously scans for the start marker and accumulates subsequent bytes until a complete frame is detected.

Once a potential frame is identified, the parser verifies its structural consistency using the payload length and integrity field. Frames that fail validation are discarded, and the parser immediately resumes searching for the next valid start marker. This process ensures rapid recovery from corrupted or misaligned data while maintaining bounded processing time.

By explicitly handling frame reconstruction at the byte level, the system maintains robustness against communication disturbances such as buffer fragmentation, partial reads, and transient transmission errors. This capability is essential for achieving reliable telemetry acquisition in non-real-time environments. The operational flow of the proposed communication framework is illustrated in Fig. 2. As shown in Fig. 2, the system operates through a structured sequence of states, including command scheduling, frame transmission, telemetry reception, validation, and parsing, ensuring consistent interaction between control and communication processes.

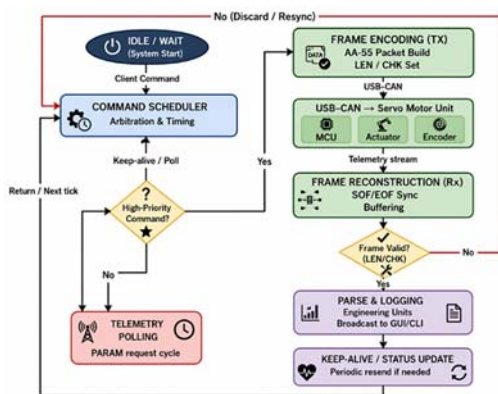


Fig. 2. Frame Reconstruction Process.

4.4. Timing Model and Communication Characteristics

The communication latency of the proposed framework can be described as a combination of

multiple processing stages across the host system, communication interface, and actuator unit. The total round-trip latency is expressed as the sum of transmission scheduling delay, interface-level buffering, bus transmission time, actuator processing delay, and reception-side parsing time. In this study, these components were not evaluated individually; instead, the total latency was analyzed as an aggregate performance measure.

$$T_{rt} = T_{tx} + T_{usb} + T_{can} + T_{servo} + T_{rx} \quad (1)$$

The actuation command is transmitted at a fixed update interval, while telemetry data are collected at a slightly slower but configurable polling rate to balance communication load and feedback availability.

$$T_c = 50 \text{ ms}, T_t \in [80 \text{ ms}, 100 \text{ ms}] \quad (2)$$

To evaluate the response behavior, the latency between command transmission and the reception of corresponding telemetry feedback is measured as follows:

$$L = t_{rx}^{telemetry} - t_{tx}^{command} \quad (3)$$

Communication load is quantified by considering the number and size of transmitted frames within a given observation window, ensuring that the bus utilization remains within a predefined safe operating range.

$$R = \frac{N_c \cdot S_c + N_t \cdot S_t}{T}, R < R_{max} \quad (4)$$

Finally, the variability in latency is characterized using statistical analysis of repeated measurements, providing insight into timing stability under non-real-time execution conditions.

$$\sigma_L = \sqrt{\frac{1}{N} \sum_{i=1}^N (L_i - \mu_L)^2} \quad (5)$$

These formulations allow systematic evaluation of communication performance and provide a basis for analyzing the impact of scheduling, buffering, and transport-layer behavior on actuator response.

5. Experimental Results and Discussion

5.1. Actuation Performance and Tracking Consistency

The effectiveness of the proposed framework was evaluated through a multi-step actuation experiment, where reference position commands were incrementally increased across a wide angular range. The objective of this evaluation was to assess the consistency between commanded inputs and measured actuator responses under the proposed communication and scheduling structure. The experimental results are

illustrated in Fig. 3, which presents the actuator position, velocity, and current profiles over time under multi-step reference inputs.

The experimental results demonstrate a high degree of agreement between the reference trajectory and the measured position values. The actuator response closely follows the commanded input throughout the entire operating range, with deviations

remaining within the resolution limits of the telemetry and encoder interface. The mean position error was approximately 0.5 deg, with a maximum deviation below 0.8 deg. This indicates that the proposed communication framework is capable of delivering commands in a stable and consistent manner without introducing significant distortion or delay.

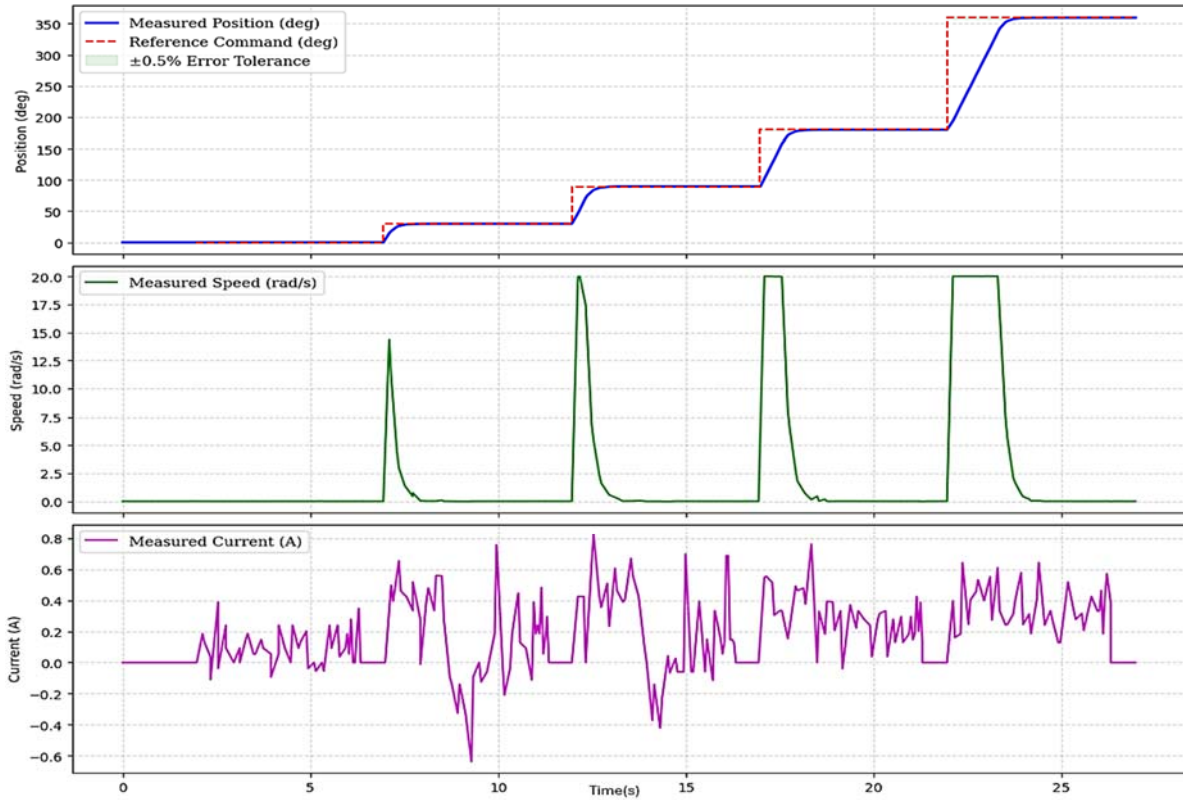


Fig. 3. Time-synchronized actuator response showing position, velocity, and current profiles under step command inputs.

In addition to positional accuracy, the dynamic behavior of the actuator was analyzed through velocity and current profiles. During each transition, the velocity increased smoothly without oscillatory behavior, while the current exhibited transient peaks corresponding to the torque required for motion initiation, followed by rapid stabilization at steady state. These observations suggest that the communication framework maintains sufficient temporal consistency to support stable actuator operation, even without direct modification of the underlying control algorithm.

5.2. Communication Stability and Timing Variability

The temporal characteristics of the proposed framework were evaluated by analyzing the intervals between consecutive telemetry receptions under periodic polling conditions. The results indicate that the telemetry reception intervals are distributed around

the intended polling range, reflecting the influence of operating system scheduling and communication buffering.

Although variability in timing is observed, the majority of telemetry frames are received within a predictable operating window. This behavior demonstrates that the system maintains practical timing consistency despite operating in a non-real-time environment. More importantly, no frame loss or data corruption was observed during the experiments, indicating that the framing structure and parsing mechanism effectively preserve communication reliability. The mean telemetry reception interval was approximately 85 ms, with most samples concentrated in the range of 70–100 ms.

5.3. Discussion and System Implications

The experimental results highlight an important characteristic of actuator communication systems in unmanned aerial platforms: the reliability of data

transport and timing structure can significantly influence system performance, even when the control algorithm itself is not modified. The proposed framework demonstrates that stable actuator behavior can be achieved by ensuring consistent command delivery and structured telemetry acquisition at the communication level.




Fig. 4. Distribution of telemetry reception intervals under periodic polling conditions, showing the variability of communication timing in a non-real-time environment.

One notable observation is that, despite the presence of timing variability in the host system, the actuator response remains stable. This suggests that communication robustness, including frame synchronization and data validation, plays a critical role in decoupling control performance from non-deterministic system behavior. In this context, the use of explicit framing and structured scheduling provides a practical means of improving system reliability without requiring a real-time operating system.

Furthermore, the modular architecture of the framework allows it to function as a diagnostic and integration tool in addition to a communication interface. By exposing transport-layer behavior such as frame timing, data reconstruction, and command-telemetry interaction, the system enables detailed analysis of communication-induced effects that are often hidden in middleware-based implementations.

However, the current approach is subject to limitations inherent to non-real-time execution environments. While the system achieves stable operation within practical timing bounds, strict real-time guarantees cannot be ensured without operating system-level support. Future work may explore integration with real-time scheduling mechanisms or hardware-assisted communication interfaces to further reduce timing variability and enhance determinism. The temporal distribution of telemetry reception intervals is further illustrated in Fig. 5. The results reveal distinct clustering patterns, with the majority of samples concentrated around the nominal operating range, while occasional deviations

appear due to system-level timing variability. These clusters reflect the combined effects of scheduling latency and communication buffering, which introduce non-uniform intervals in telemetry acquisition. Despite these variations, the overall distribution remains stable, indicating that the proposed framework maintains consistent communication behavior within practical operating bounds.




Fig. 5. Temporal distribution of telemetry reception intervals over successive samples, illustrating timing variability and clustering behavior under non-real-time communication conditions.

5.4. Comparative Evaluation

To validate the effectiveness of the proposed framework, a comparative analysis was conducted against conventional actuator communication approaches, including PWM-based control and ROS2 middleware-based communication.

PWM-based control provides low-latency signal transmission but lacks structured telemetry feedback and communication-level validation, limiting its suitability for applications requiring high observability and reliability. In contrast, ROS2-based middleware enables modular system integration and structured message exchange, but its performance is strongly affected by operating system scheduling and communication overhead, resulting in increased latency and timing variability.

The proposed USB-CAN framework achieves a balanced performance by combining structured communication with explicit timing control. Experimental observations indicate that the proposed method maintains lower jitter and higher reliability compared to ROS2-based approaches, while providing significantly improved observability over PWM-based systems. Table 2 summarizes the quantitative comparison of the evaluated communication methods. The results demonstrate that the proposed framework provides a practical trade-off between latency, reliability, and system transparency, making it suitable for UAV actuator systems operating in non-real-time

environments. The values in Table 2 are based on direct measurements for the proposed framework and representative estimates derived from system

characteristics and prior studies for conventional approaches.

Table 2. Quantitative comparison of actuator communication methods based on experimental observations and system characteristics.

Metric	PWM (Pulse Signal)	ROS2 Middleware (DDS-Based)	USB-CAN (Proposed) (AA-55 / C4/C6 Frame)
Latency (Round-Trip) [ms]	5 – 10 Low latency due to simple pulse transmission	20 – 50 Affected by middleware overhead and OS scheduling	10 – 20 Moderate latency with structured framing and scheduling
Jitter (Std. Dev.) [ms]	Low Stable signal timing with minimal variation	High Variability due to non-real-time infrastructure	Low Bounded variability with dual-loop scheduling
Reliability	Medium No error detection / No feedback validation	Medium Reliable transport, but timing variability may cause loss	High Frame validation (AA-55) and checksum ensure high reliability
Observability	Low No telemetry feedback (open-loop)	Medium Telemetry available, but debugging affected by complexity	High Structured telemetry with explicit parsing and logging
Implementation Complexity	Low Simple hardware and software design	High Requires complex middleware configuration	Medium Custom framing and scheduling, but modular and lightweight
Bus Utilization (Measured / Estimated)	N/A (Analog pulse signal)	~40 – 70 % Depends on message rate and system load	< 35 % Measured under experimental conditions

* Values are measured or estimated based on experimental observations in a non-real-time operating environment.

6. Conclusion

This study presented a communication-oriented servo actuation framework designed for unmanned aerial systems using a universal serial bus to controller area network interface. The proposed approach focused on improving the reliability of command transmission and telemetry acquisition rather than modifying the underlying control algorithm. By introducing a structured dual-loop scheduling mechanism and a custom byte-level framing scheme, the framework enables consistent data exchange and explicit handling of frame synchronization and validation under non-real-time operating conditions.

Experimental results demonstrated that the proposed system achieves high consistency between commanded and measured actuator responses, with deviations remaining within the resolution limits of the telemetry and encoder interface. In addition, the communication framework maintained stable telemetry acquisition despite timing variability introduced by operating system scheduling and interface buffering. These results confirm that reliable actuator behavior can be achieved through robust communication design even in environments where strict real-time guarantees are not available.

The proposed framework provides a practical and modular solution for actuator communication in unmanned aerial platforms, offering improved transparency in transport-layer behavior and facilitating integration with higher-level systems.

Future work will focus on extending the framework toward real-time operating environments and multi-actuator configurations to further enhance timing determinism and system scalability.

Acknowledgement

The current research was supported by the Space Challenge Program funded by the Korea government (Korea AeroSpace Administration, KASA) (RS-2022-NR067041).

This work is supported by the Korea Agency for Infrastructure Technology Advancement (KAIA) grant funded by the Ministry of Land, Infrastructure and Transport (Grant RS-2024-00408551).

References


- [1]. C. Anastasopoulos, G. Tsourveloudis, K. Dalamagkidis, Design of a real-time test bench for UAV servo actuators, in *Proceedings of the International Conference on Unmanned Aircraft Systems (ICUAS'18)*, 2018, pp. 1224-1231.
- [2]. J. Sharma, R. Kumar, A. Nayyar, Robust PID control of single-axis gimbal actuator via UAV platform, *IFAC-PapersOnLine*, Vol. 53, Issue 1, 2020, pp. 113-118.
- [3]. E. M. Coates, D. Reinhardt, K. Gryte, T. A. Johansen, Toward nonlinear flight control for fixed-wing UAVs: System architecture, field experiments, and lessons

learned, in *Proceedings of the International Conference on Unmanned Aircraft Systems (ICUAS'22)*, 2022, pp. 715-725.

- [4]. J. Lee, S. Cho, Design and implementation of a servo motor control library for drone control systems, in *Proceedings of the 2nd International Conference on Drones and Unmanned Systems (DAUS'26)*, 2026, pp. 53-58.
- [5]. T. Kronauer, J. Pohlmann, M. Matthé, T. Smejkal, et al., Latency analysis of ROS2 multi-node systems, in *Proceedings of the IEEE International Conference*

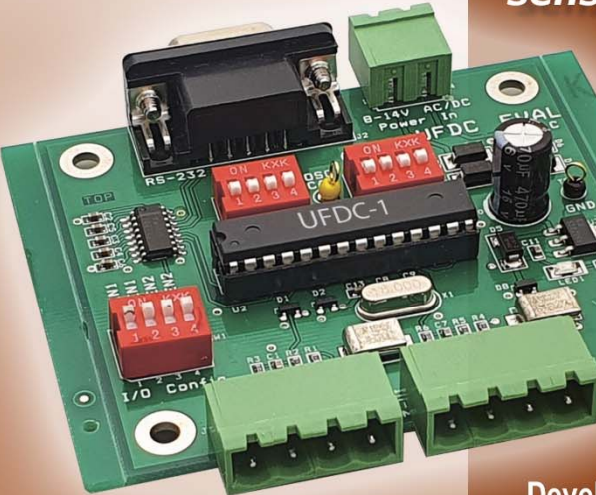

on Multisensor Fusion and Integration for Intelligent Systems (MFI'21), 2021, pp. 1-7.

- [6]. Y. Ye, Z. Nie, X. Liu, X. Wu, et al., ROS2 real-time performance optimization and evaluation, *Chinese Journal of Mechanical Engineering*, Vol. 36, 2023, 144.
- [7]. OpenCyphal Development Team, Cyphal specification v1.0, 2025, <https://specification.opencyphal.org>



Easy and quick

sensors systems development

Development Board UFDC-1/UFDC-1M-16

- 16 measuring modes, 2 channels for frequency measurements
- Frequency range from 0.05 Hz up to 7.5 MHz (120 MHz)
- Programmable accuracy from 1 % up to 0.001 %
- RS232 (USB optional)

sales@sensorsportal.com
http://www.sensorsportal.com/HTML/E-SHOP/PRODUCTS_4/Evaluation_board.htm



A Security-Aware Edge Artificial Intelligence Framework for Robust Autonomous Decision-Making in Unmanned Aerial Systems

Sérgio SILVA

Department of Communication Sciences and Information Technologies,
University of Maia, 4475-690 Maia, Portugal
E-mail: D012196@umaia.pt

Received: 26 Dec. 2025 / Revised: 31 March 2026 / Accepted: 20 April 2026 / Published: 28 April 2026

Abstract: This article presents a security-aware Edge Artificial Intelligence framework designed to enhance autonomous decision-making in Unmanned Aerial Systems. As aerial vehicles increasingly rely on onboard artificial intelligence to interpret sensor data, plan flight trajectories, and respond to environmental conditions, they are exposed to cyber-physical threats such as data manipulation, Global Navigation Satellite System (GNSS) spoofing, and adversarial visual perturbations. The proposed framework integrates lightweight deep learning models optimized for embedded processors with a multilayer cybersecurity architecture that includes real-time integrity verification of sensor streams, adversarial robustness modules, and secure decision-validation routines. Experimental evaluation on a quadrotor platform demonstrates that the system preserves autonomy performance while significantly reducing the risk of unsafe commands under adversarial conditions, achieving over an 85 % reduction in unsafe actions with minimal false alarms and negligible impact on mission continuity. These findings highlight the importance of combining Edge Artificial Intelligence with embedded security mechanisms to ensure the resilience, safety, and reliability of autonomous unmanned systems operating in contested or communication-limited environments.

Keywords: Edge artificial intelligence, Unmanned aerial systems, UAV cybersecurity, Autonomous decision-making, Embedded AI, Adversarial robustness, Sensor integrity monitoring.

1. Introduction

Autonomous Unmanned Aerial Systems (UAS) have advanced rapidly due to the integration of onboard Artificial Intelligence (AI), enabling complex missions with limited human intervention. These systems rely on digital sensors, communication links, and machine learning models to interpret environmental data, plan trajectories, and execute decisions in real-time. However, this increasing reliance on digital intelligence exposes unmanned systems to a range of cyber-physical threats, including data manipulation, GNSS spoofing, and adversarial perturbations in visual sensors. Such attacks can compromise decision-making, reduce operational safety, and threaten mission success.

Recent research has increasingly focused on improving the security and reliability of autonomous Unmanned Aerial Systems (UAS) through the integration of AI and advanced anomaly detection techniques. Several studies have explored the use of Machine Learning models to detect cyber-physical threats affecting aerial platforms. For example, AI-based frameworks have been proposed to detect cyber-physical attacks in Internet-of-Things enabled drones by monitoring system behavior and identifying abnormal patterns in sensor data [1].

Other works have specifically addressed GNSS spoofing attacks, which represent one of the most critical threats to autonomous navigation. Machine learning and adversarial learning techniques have been investigated to identify spoofing patterns and mitigate

navigation errors in Unmanned Aerial Vehicles (UAV) [2, 3].

Recent approaches also explore ensemble learning and deep learning methods for detecting navigation signal manipulation and improving the robustness of positioning systems in aerial platforms [4].

In addition, research has been conducted on explainable reinforcement learning methods capable of detecting adversarial attacks affecting guidance and path planning modules [5].

Experimental studies have also demonstrated the feasibility of detecting spoofing attacks through sensor-fusion analysis and recovery mechanisms based on consistency verification between inertial and navigation data [6-8].

Recent advances in autonomous aerial systems have highlighted the importance of integrating AI with cybersecurity mechanisms to ensure safe operation in dynamic environments [9].

Several studies have investigated the use of deep learning models for real-time perception and navigation in unmanned aerial vehicles, demonstrating significant improvements in obstacle detection and environment understanding [10, 11].

However, the integration of such perception systems with security-aware mechanisms remains a relatively unexplored research area. In particular, the challenge of ensuring trustworthy autonomous decision-making under adversarial conditions has motivated the development of embedded security frameworks capable of detecting sensor anomalies, verifying data integrity, and validating navigation commands before execution [12, 13].

Despite these advances, many existing solutions focus primarily on individual attack detection mechanisms and often rely on centralized processing or external monitoring systems. This can limit their applicability in real-time autonomous platforms operating with constrained computational resources or in communication-limited environments [14, 15]. Therefore, there remains a need for integrated frameworks that combine embedded AI perception with real-time cybersecurity mechanisms capable of validating both sensor data and autonomous decisions directly on the onboard computing platform [9, 16].

1.1. Extension over Previous DAUS 2026 Conference Paper

This article extends previous work presented at the DAUS 2026 conference [17]. While the conference version introduced the initial concept of a secure Edge AI framework, the present manuscript provides significant extensions. These include enhanced real-time security layers, a more detailed and formalized system architecture, and expanded experimental evaluation under multiple adversarial scenarios. In addition, this work introduces improved implementation details and formally defined evaluation metrics, strengthening reproducibility and scientific rigor.

The key differences between the conference and journal versions are summarized in Table 1.

Table 1. Conference versus Journal versions.

Component	DAUS 2026	Journal article
Architecture	Conceptual design	Full multi-layer implementation
EKF Monitoring	Basic description	Formal EKF with statistical detection
Perception Model	Preliminary YOLOv5	Optimized and embedded deployment
Experiments	Limited validation	Multiple adversarial scenarios
Evaluation Metrics	Partial	Comprehensive and formally defined
Reproducibility	Limited details	Full implementation detail

The main contributions of this work can be summarized as follows:

- A security-aware Edge AI architecture integrating perception, integrity monitoring, and decision validation for autonomous UAS;
- A real-time sensor integrity monitoring approach based on sensor fusion and anomaly detection techniques;
- A secure decision validation mechanism that prevents unsafe or compromised commands from being executed;
- An embedded implementation optimized for resource-constrained UAV platforms;
- Experimental validation under multiple cyber-physical attack scenarios, demonstrating improved system resilience and safety.

The main novelty of this work lies in the integration of lightweight perception models, sensor integrity monitoring, and secure decision validation within a unified Edge AI architecture capable of operating fully onboard resource-constrained UAV platforms.

The remainder of this article is organized as follows. Section 2 presents the proposed methodology. Section 3 describes the system architecture and implementation. Section 4 provides experimental results and evaluation. Section 5 discusses the findings and limitations, and Section 6 concludes the article and outlines future work.

2. Methodology

The proposed framework integrates secure Edge AI capabilities into the autonomous decision-making pipeline of UAS. As illustrated in Fig. 1, the methodology consists of three main components: (i) onboard perception and inference, (ii) cybersecurity and data integrity verification, and (iii) secure decision validation. All components are

optimized to operate in real-time on embedded hardware with limited computational resources, ensuring reliable and safe autonomy even under adversarial conditions.

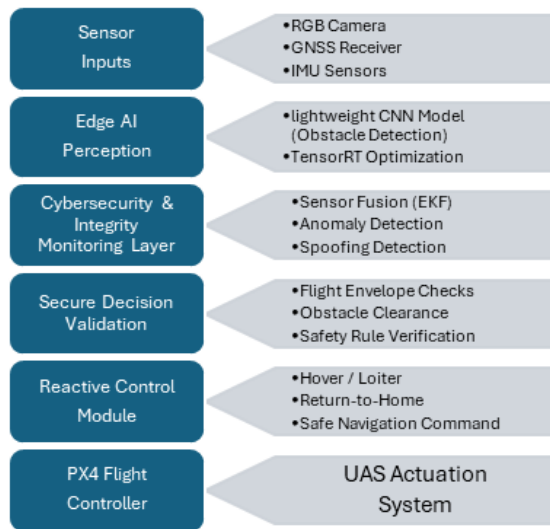


Fig. 1. Architecture of secure Edge AI framework for autonomous UAS.

Fig. 1 illustrates the overall architecture of the proposed secure Edge AI framework. Sensor data from the onboard camera, global navigation satellite receiver, and inertial measurement unit are first processed by the perception and sensor fusion modules.

The perception module performs obstacle detection using a lightweight Convolutional Neural Network (CNN), while the integrity monitoring layer evaluates the consistency of navigation and inertial data. The perception layer processes visual information acquired from the onboard camera using and extracts environmental features and identifies potential obstacles in real-time.

Detected anomalies are forwarded to the decision validation module, which enforces safety constraints before navigation commands are transmitted to the flight controller.

The sensor fusion and integrity monitoring layer combines data from the GNSS receiver and inertial measurement unit through an Extended Kalman Filter (EKF). By evaluating the innovation residual between predicted and observed measurements, the system can identify inconsistencies that may indicate sensor spoofing or data manipulation attacks.

The final layer performs secure decision validation by enforcing operational safety constraints on the navigation commands generated by the autonomous control system. If anomalies are detected or safety conditions are violated, the system triggers fallback behaviors such as hover or return-to-home maneuvers.

This layered architecture provides multiple levels of protection against cyber-physical threats while

maintaining real-time operation on resource-constrained embedded platforms.

2.1. Onboard Perception and Edge AI Inference

The onboard perception module performs obstacle detection and free-space assessment from monocular RGB imagery. A compressed CNN, based on Nano YOLOv5 and optimized for embedded deployment, processes 416×416 images to output obstacle bounding boxes with confidence scores. Training combines the VisDrone dataset with in-house outdoor flight data using augmentation techniques such as brightness variation, Gaussian noise, and motion blur.

Real-time performance on an NVIDIA Jetson Nano is achieved through post-training INT8 quantization, structured channel pruning, and TensorRT optimization. The model achieves a mean average precision (mAP@0.5) of 71.4 %, precision of 0.76, recall of 0.69, and an inference rate of 18–22 frames per second, depending on scene complexity. These results demonstrate that lightweight deep learning models can provide accurate perception while meeting embedded hardware constraints.

The training dataset combines publicly available aerial imagery from the VisDrone dataset with additional flight recordings captured during outdoor test missions. The combined dataset includes urban, semi-urban, and natural environments, providing diverse visual conditions for training the perception model. Data augmentation techniques were applied to simulate variations in illumination, motion blur, and sensor noise. These augmentations improve the robustness of the perception model when operating under challenging environmental conditions commonly encountered during real-world UAV deployments.

To improve reproducibility, the implementation details of the perception model are specified as follows. The Nano YOLOv5 model corresponds to a compressed variant of YOLOv5n adapted for embedded deployment. The training dataset consists of approximately 18,500 annotated images, combining the VisDrone dataset (70 %) and in-house UAV flight data (30 %). The dataset was divided into training (70 %), validation (20 %), and test (10 %) subsets.

Training was performed for 150 epochs using the Adam optimizer with an initial learning rate of 0.001 and cosine decay scheduling. A batch size of 16 was used during training. Data augmentation included brightness variation, Gaussian noise, motion blur, and horizontal flipping.

During inference, a confidence threshold of 0.4 and a non-maximum suppression threshold of 0.5 were applied. Deployment on the NVIDIA Jetson Nano was achieved using TensorRT with INT8 quantization calibrated on a subset of 1,000 representative images. Structured pruning reduced the model size by

approximately 35 % while preserving detection accuracy.

These details enable reproducibility and allow fair comparison with other embedded perception approaches.

2.2. Cybersecurity and Data Integrity Layer

A dedicated integrity-monitoring layer evaluates the consistency of sensor data from GNSS, inertial measurement units, and vision-based estimates using an EKF. Anomaly detection relies on innovation residuals, triggering alerts when thresholds are exceeded. Monitored features include position drift, velocity inconsistencies, inertial sensor bias variation, and disagreement between vision and inertial measurements.

The EKF used for sensor fusion estimates the system state vector, including position, velocity, and orientation of the aerial vehicle. The filter continuously compares predicted and measured sensor values through innovation residuals. When the residual exceeds a predefined statistical threshold, the system identifies a potential anomaly in the sensor measurements. This approach enables the detection of navigation inconsistencies caused by spoofing attacks, sensor faults, or malicious data manipulation.

This layer ensures that the system can detect potential cyber-physical attacks, including sensor spoofing and data injection, within minimal latency. Fig. 1 illustrates the integration of the perception and integrity-monitoring modules in the overall system architecture.

2.3. Secure Decision-Making and Validation Layer

Autonomous navigation commands generated by the onboard AI are validated by a safety filter that enforces flight envelope constraints, obstacle clearance, and integrity alerts. Commands associated with detected anomalies are rejected, triggering a reactive control module that executes safe fallback actions such as hover, loiter, or return-to-home maneuvers.

This approach prevents unsafe actions while maintaining mission continuity. Table 2 summarizes the system performance under various adversarial scenarios, highlighting reductions in unsafe commands, false alarm rates, and reaction latency.

The interaction between the perception, integrity monitoring, and decision validation modules forms a layered defense architecture designed to improve system robustness. The perception module provides environmental awareness through visual analysis, while the integrity monitoring layer evaluates the reliability of sensor measurements and detects inconsistencies that may indicate cyber-physical attacks. The decision validation module enforces a

formally defined safety envelope to prevent unsafe navigation commands. The constraints include maximum linear velocity (5 m/s), maximum angular rate (120 °/s), and minimum obstacle clearance distance (1.5 m). In addition, navigation commands are rejected when an integrity alert is active.

A command is classified as unsafe if it results in a predicted trajectory that violates any of these constraints or introduces a collision risk. When such a condition is detected, the command is discarded and a fallback control strategy is activated.

The fallback mechanism follows a hierarchical structure: (i) hover stabilization as default response, (ii) loiter mode when localization remains stable, and (iii) return-to-home if anomalies persist for more than 2 seconds. This ensures safe operation while minimizing mission disruption.

This formalization transforms the decision validation layer into a deterministic safety filter, strengthening the reliability of autonomous decision-making under adversarial conditions.

2.4. Sensor Fusion and Anomaly Detection Model

The integrity monitoring layer relies on an EKF to estimate the system state using measurements from the inertial measurement unit, visual perception module, and GNSS. The state vector can be represented as:

$$\mathbf{x}_k = [\mathbf{p}_k \ \mathbf{v}_k \ \theta_k],$$

where \mathbf{p}_k represents the position, \mathbf{v}_k the velocity, and θ_k the orientation of the aerial vehicle at time step k .

The prediction stage follows the standard EKF formulation:

$$\hat{\mathbf{x}}_k|_{k-1} = f(\hat{\mathbf{x}}_{k-1}, \mathbf{u}_k),$$

while the update stage incorporates sensor observations:

$$\hat{\mathbf{x}}_k = \hat{\mathbf{x}}_k|_{k-1} + K_k(\mathbf{z}_k - h(\hat{\mathbf{x}}_k|_{k-1})),$$

where K_k is the Kalman gain and \mathbf{z}_k represents the measurement vector.

Anomaly detection is performed using the normalized innovation squared (NIS), which provides a statistically grounded measure of consistency between predicted and observed measurements. The innovation residual is defined as:

$$\mathbf{r}_k = \mathbf{z}_k - h(\hat{\mathbf{x}}_k|_{k-1})$$

The anomaly score is computed as:

$$\epsilon_k = \mathbf{r}_k^T \mathbf{S}_k^{-1} \mathbf{r}_k,$$

where S_k represents the innovation covariance matrix. An anomaly is triggered when:

$$\varepsilon_k > \chi^2(\alpha, d),$$

where $\chi^2(\alpha, d)$ is the chi-square threshold corresponding to confidence level α and d degrees of freedom. In this work, a confidence level of $\alpha = 0.99$ was selected to balance detection sensitivity and false positive rate.

The process and measurement noise covariance matrices were empirically calibrated using nominal flight data. Sensitivity analysis showed that variations of $\pm 10\%$ in the threshold resulted in less than 3% variation in detection performance, indicating robustness of the detection mechanism.

2.5. System Architecture Overview

The overall system architecture follows the layered design introduced in Section 2 and is implemented on an embedded UAV platform. The framework integrates three main components: a perception module, an integrity monitoring module, and a decision validation module.

Rather than repeating the methodological details, this section focuses on the practical integration of these components within the onboard system. The perception module processes visual data in real-time using the optimized Nano YOLOv5 model. The integrity monitoring module evaluates sensor consistency through EKF-based estimation, while the decision validation layer filters navigation commands based on predefined safety constraints.

All modules operate sequentially within the onboard computing pipeline, ensuring that perception outputs and state estimates are continuously verified before control actions are executed. This integrated design enables real-time operation under resource constraints while maintaining robustness against adversarial conditions.

Implementation details, including hardware configuration and optimization techniques, are provided in the following subsection.

2.6. Threat Model

The security analysis of the proposed framework considers a cyber-physical adversary capable of manipulating sensor data or navigation signals during autonomous flight operations. The attacker may attempt to influence the decision-making pipeline by injecting corrupted measurements into the navigation or perception modules.

Three primary attack vectors are considered in this study:

- Navigation spoofing attacks, where the adversary manipulates GNSS signals to

introduce false position estimates and alter the flight trajectory;

- Adversarial perception attacks, where carefully crafted perturbations are introduced into camera images to mislead the deep learning perception model;
- Sensor data injection, where malicious values are inserted into inertial or navigation sensor streams at the middleware level.

The attacker is assumed to have the capability to interfere with external sensor inputs but does not have direct access to the internal flight controller firmware or onboard computing hardware. This assumption reflects realistic threat scenarios where adversaries attempt to manipulate external signals or communication interfaces rather than compromising the entire onboard system.

Under this threat model, the proposed security-aware Edge AI framework focuses on detecting inconsistencies between multiple sensor sources and preventing compromised data from influencing the autonomous control pipeline.

2.7. Algorithm: Secure Edge AI Decision Pipeline

Input: Sensor data streams (camera, GNSS, IMU)

Output: Validated navigation command:

- Acquire sensor data from camera, GNSS receiver, and IMU;
- Perform visual perception using compressed CNN model;
- Estimate vehicle state using EKF-based sensor fusion;
- Compute innovation residual between predicted and observed measurements;
- If residual exceeds anomaly threshold
Trigger integrity alert;
- Apply decision validation filter to navigation commands;
- If command violates safety constraints or integrity alert is active
Execute fallback maneuver (hover or return-to-home);
- Otherwise
Transmit validated command to flight controller.

3. Experimental Setup and Results

This section presents the experimental evaluation of the proposed secure Edge AI framework. The objective of the experiments is to assess the capability of the system to detect and mitigate cyber-physical threats while maintaining real-time autonomous operation on embedded hardware.

The evaluation focuses on three representative attack scenarios affecting autonomous UAS, including navigation signal spoofing, visual adversarial perturbations, and sensor data manipulation.

Detection accuracy, false positive rate, latency, and inference speed were measured to assess system performance.

All flight experiments were conducted in controlled outdoor environments following standard UAV safety procedures.

3.1. Experimental Platform

Experiments were conducted on a quadrotor UAS equipped with an embedded computing platform and multiple onboard sensors. The embedded computing platform used in the experiments provides a balance between computational performance and energy efficiency, which is critical for small aerial vehicles. The NVIDIA Jetson Nano offers a GPU-accelerated architecture capable of executing deep learning inference while maintaining relatively low power consumption. This makes it suitable for real-time Edge AI applications in autonomous unmanned systems.

The PX4 autopilot stack was used to control the flight dynamics of the quadrotor platform and to interface with the onboard sensors. The quadrotor platform used in the experiments has a diagonal wheelbase of 450 mm and a total takeoff weight of approximately 1.8 kg. The onboard computing module operates under Ubuntu Linux and runs the Robot Operating System (ROS) middleware for sensor integration and communication between perception, navigation, and security modules. Sensor data streams are synchronized through ROS topics, allowing real-time access to navigation and perception information by the integrity monitoring processes.

The Edge AI inference module was implemented using TensorRT optimization to achieve real-time performance on the embedded hardware. The integrity monitoring and security verification processes were executed as companion processes on the same platform, allowing synchronized access to sensor data

and navigation commands. The complete system architecture is illustrated in Fig. 1.

3.2. Test Scenarios

Three representative cyber-physical threat scenarios were implemented to assess robustness during controlled flight experiments.

The first scenario simulates a global navigation satellite spoofing attack by gradually injecting position drift into the navigation data stream. This attack aims to cause navigation errors and potentially unsafe trajectory deviations.

The second scenario introduces adversarial perturbations into the visual perception pipeline. These perturbations are generated using a fast gradient sign method applied to camera frames to evaluate the robustness of the perception model against manipulated visual inputs.

The third scenario simulates sensor data injection by modifying inertial measurement unit readings at the middleware level. This test evaluates the ability of the integrity monitoring layer to detect inconsistencies between inertial, visual, and navigation measurements.

Ground truth reference data were obtained from unaltered sensor logs and reference flights recorded under normal operating conditions.

3.3. Experimental Results

Each scenario was evaluated through twenty flight trials with an average duration of approximately five minutes. Detection accuracy, false positive rate, detection latency, and real-time inference performance were measured for each threat type. Table 2 summarizes the results obtained.

The results demonstrate that the proposed security-aware architecture maintains high detection accuracy across different threat scenarios while preserving real-time system performance. The relatively low detection latency enables the activation of safety mechanisms before compromised sensor data can significantly affect flight control decisions.

Table 2. Performance of the secure Edge AI framework under adversarial scenarios.

Threat Scenario	Detection Rate (%)	False Positive Rate (%)	Detection Latency (ms)	Inference Speed (FPS)
GNSS Spoofing	92 ± 3	4.1	180	20
Visual Perturbation	88 ± 4	5.3	210	18
Sensor Data Injection	90 ± 2	4.7	165	22
Average	90	4.7	185	20

Detection latency remains below 210 ms across all evaluated scenarios, allowing the security mechanisms to respond quickly and prevent unsafe navigation commands. In addition, the perception module maintains real-time inference speeds between 18 and 22 frames per second, demonstrating that the integration of cybersecurity mechanisms does not significantly degrade system performance.

3.4. Ablation Study

An ablation study was also conducted to evaluate the contribution of individual system components. The comparison included four configurations: a baseline system without security mechanisms, a system using only the integrity monitoring layer, a system using

only the safety validation filter, and the full proposed architecture. The results are presented in Table 3.

The ablation study highlights the importance of combining integrity monitoring and decision validation within a unified architecture. While individual modules provide moderate improvements in detection performance, the complete system significantly reduces unsafe commands and improves operational resilience. This demonstrates that layered security mechanisms are essential for protecting autonomous aerial systems operating in adversarial environments.

The complete system reduced unsafe commands by more than 85 % compared with the baseline configuration while maintaining a mission performance degradation of less than 3 % during normal operation.

The results demonstrate that integrating integrity monitoring and secure decision validation with Edge AI significantly improves the resilience of autonomous UAS operating in adversarial environments.

To further evaluate the effectiveness of the proposed framework, its performance was compared with several recent approaches addressing security and attack detection in autonomous UAS.

3.5. Comparative Evaluation

The comparison focuses on detection accuracy, response latency, and suitability for embedded real-time deployment. The results of this comparison are summarized in Table 4.

Table 3. Ablation Study of Security Components.

System Configuration	Detection Rate (%)	Unsafe Commands	Mission Degradation (%)
Baseline (no security)	42	High	0
Integrity Monitoring only	78	Medium	1.2
Safety Validation only	73	Medium	1.5
Proposed Full Architecture	92	Low	2.8

Table 4. Comparative evaluation of the proposed framework with existing UAV security approaches.

Method	Attack Detection Rate (%)	Detection Latency (ms)	Real-time Embedded Support	Security Scope
ML-based GNSS spoofing detection	85	250	Limited	Navigation attacks
Deep learning UAV protection	87	230	Partial	GNSS spoofing
Ensemble spoofing detection	89	240	No	Navigation attacks
Reinforcement learning attack detection	90	300	No	Guidance attacks
Proposed Secure Edge AI Framework	92	180	Yes	Multi-sensor cyber-physical threats

3.6. Comparative Analysis with Existing Approaches

The comparative results presented in Table 4 highlight several significant differences between the proposed framework and existing approaches for securing autonomous UAS. While previous studies have demonstrated promising results in detecting specific attack vectors such as GNSS spoofing or navigation manipulation, most of these methods focus on a single subsystem of the UAV architecture.

For instance, machine learning-based spoofing detection approaches typically analyze navigation signals in isolation, without considering cross-validation with other onboard sensors. Similarly, deep learning protection mechanisms are often designed specifically for visual perception pipelines and may not address inconsistencies originating from navigation or inertial measurements.

As a result, these solutions may fail to detect coordinated cyber-physical attacks that simultaneously affect multiple sensor streams.

By contrast, this approach uses a multi-layer architecture that unifies perception, integrity monitoring, and decision validation within the Edge AI pipeline. By integrating information from multiple onboard sensors through an EKF-based fusion process, the system can identify inconsistencies between navigation, inertial, and visual measurements. This cross-sensor validation significantly improves the reliability of anomaly detection compared with approaches relying on a single data source.

Another important distinction concerns computational feasibility for embedded deployment. Several existing solutions rely on computationally intensive deep learning models or require centralized processing infrastructures for anomaly detection. Such

approaches may introduce significant communication latency or may not be suitable for small UAV platforms with limited processing resources. To overcome this limitation, the system leverages optimized neural networks and integrity monitoring algorithms suitable for embedded platforms like the NVIDIA Jetson Nano.

From a performance perspective, the proposed framework achieves a detection rate of approximately 92 % with an average response latency of 180 ms, outperforming or matching the detection capabilities of several state-of-the-art approaches while maintaining full compatibility with real-time embedded operation. The reduced detection latency is particularly important for autonomous aerial systems, where delayed responses to cyber-physical attacks may lead to unsafe flight behaviors.

Furthermore, the proposed decision validation layer introduces an additional safety barrier that is not commonly present in existing detection frameworks. While many approaches focus exclusively on identifying anomalies, the architecture presented in this study actively prevents compromised data from influencing flight control commands by enforcing safety constraints and triggering fallback control behaviors.

Overall, the comparative analysis indicates that integrating perception, cybersecurity monitoring, and decision validation within a unified Edge AI architecture provides a more comprehensive and resilient protection strategy for autonomous UAS. This integrated design enables reliable operation in adversarial environments while maintaining the computational efficiency required for deployment on resource-constrained aerial platforms.

3.7. Evaluation Metrics

To ensure clarity and reproducibility, the evaluation metrics are formally defined as follows.

Detection rate represents the percentage of attack instances correctly identified by the system. False positive rate indicates the proportion of normal operating conditions incorrectly classified as anomalies.

Detection latency corresponds to the time elapsed between the onset of an attack and its detection, which is a critical parameter for real-time autonomous systems.

An unsafe command is defined as any navigation command that leads to a predicted trajectory violating safety constraints, including obstacle proximity below 1.5 m, excessive velocity, or instability under inconsistent sensor data.

Mission degradation is defined as the relative increase in mission completion time compared to baseline operation:

$$\text{Degradation}(\%) = \frac{T_{\text{attack}} - T_{\text{baseline}}}{T_{\text{baseline}}} \times 100$$

All metrics were computed over 20 flight trials per scenario, each lasting approximately five minutes, ensuring statistical consistency.

4. Discussion

The experimental results demonstrate that integrating cybersecurity mechanisms directly within the Edge AI processing pipeline significantly enhances the operational resilience of autonomous UAS. The system effectively detects and mitigates cyber-physical threats while maintaining real-time performance on embedded hardware, highlighting the feasibility of onboard security-aware intelligence.

A key observation is the effectiveness of the integrity monitoring layer in detecting inconsistencies among navigation, inertial, and vision-based estimates with low latency. This capability is particularly important in scenarios where GNSS signals are unreliable or intentionally manipulated. Early detection of anomalies prevents corrupted data from propagating through the decision pipeline, thereby reducing the risk of unsafe system behavior.

The results also underline the importance of decision-level safety enforcement. While anomaly detection identifies suspicious inputs, the validation mechanism ensures that unsafe commands are not executed. This complementary interaction between detection and validation contributes to a significant reduction in unsafe actions, while fallback strategies such as hover stabilization and return-to-home preserve mission continuity.

Another important outcome concerns the balance between security and operational efficiency. The results indicate that improved resilience does not significantly degrade mission performance, suggesting that security-aware processing can be integrated without compromising real-time constraints. This is particularly relevant for embedded platforms, where computational resources are limited.

From a broader perspective, the findings highlight the potential of integrating perception, integrity monitoring, and decision validation within a unified onboard framework. Such integration enables faster response times compared to centralized approaches and improves robustness against multiple attack vectors. This is especially relevant for autonomous systems operating in communication-constrained or adversarial environments.

Despite these advantages, several limitations remain. The current implementation focuses on sensor integrity and adversarial perception attacks, while other threats such as communication link interference or coordinated multi-sensor attacks are not explicitly addressed. Furthermore, the system relies on predefined thresholds and models, which may require adaptation in highly dynamic environments. Future work will explore adaptive and learning-based security mechanisms capable of identifying previously unseen attack patterns.

From an application standpoint, the proposed approach is particularly relevant for mission-critical operations such as infrastructure inspection, environmental monitoring, and emergency response. In these scenarios, reliable onboard detection of sensor anomalies and prevention of unsafe navigation decisions are essential for ensuring safe and continuous operation, especially when communication with ground control is limited.

Overall, the results support the integration of security-aware mechanisms directly within the onboard decision pipeline as an effective strategy for improving the reliability, safety, and autonomy of next-generation aerial systems.

5. Conclusions

The experimental evaluation demonstrates that integrating cybersecurity monitoring directly within the Edge AI pipeline significantly improves the resilience of autonomous UAV systems against cyber-physical threats.

This article presented a security-aware Edge AI framework aimed at improving the resilience of autonomous decision-making in UAS. These findings indicate that security-aware Edge AI architectures represent a promising direction for enabling safe and trustworthy autonomous aerial systems in complex operational environments.

Experimental evaluation on a quadrotor platform demonstrated that the system could detect and mitigate several cyber-physical threats, including navigation spoofing, visual adversarial perturbations, and sensor data manipulation.

The results obtained in this study suggest that integrating security mechanisms directly within Edge AI architectures can significantly improve the reliability of autonomous systems operating in complex environments. Such approaches will play a crucial role in enabling safe large-scale deployment of unmanned aerial platforms in civilian, industrial, and emergency response applications.

The study highlights the importance of incorporating cybersecurity mechanisms directly within the onboard AI pipeline of autonomous aerial systems. Such integration enables safer deployment of unmanned platforms in environments where communication links may be unreliable, or adversarial conditions may be present.

Another important aspect of the proposed framework concerns computational efficiency. The perception model was optimized through model compression techniques including quantization and structured pruning, enabling efficient execution on resource-constrained hardware. Experimental results demonstrate that the integrated architecture maintains real-time inference speeds while simultaneously executing cybersecurity monitoring processes. This confirms the feasibility of deploying security-aware AI systems directly on embedded aerial platforms

without requiring high-performance ground infrastructure.

Future research will focus on addressing additional threat vectors, including communication-level attacks and cooperative multi-drone scenarios. Further improvements will also explore adaptive security mechanisms and more advanced perception models capable of improving robustness in complex operational environments.

6. Limitations and Future Work

Although the proposed framework demonstrates promising results for improving the resilience of autonomous UAS, several limitations should be acknowledged.

First, the current implementation focuses primarily on attacks affecting sensor integrity and visual perception. Other types of cyber-physical threats, such as communication link interference, coordinated multi-drone attacks, or malicious firmware modifications, were not addressed in the present study. Developing capabilities to detect and mitigate these threats is a key direction for future research.

Second, the evaluation was conducted on a single quadrotor platform equipped with a specific embedded computing device. While the results demonstrate the feasibility of secure Edge AI on resource-constrained hardware, further experiments on different aerial platforms and processing architectures would provide a more comprehensive assessment of system scalability and robustness.

Another promising direction for future research involves the integration of cooperative security mechanisms in multi-drone systems. In such environments, multiple UAV may share sensor information and collaboratively detect cyber-physical anomalies affecting the swarm. Developing distributed security-aware decision frameworks for cooperative aerial systems could further enhance resilience and mission reliability in complex operational scenarios.

Future research will focus on extending cybersecurity coverage, enhancing adaptive detection, and validating the approach in larger UAV fleets and real operational environments.

The complete system operates within the computational constraints of a low-power embedded platform, demonstrating that security-aware Edge AI can be effectively deployed in small UAV systems without requiring high-performance computing resources.


References

- [1]. A. Hakeem, M. Sabir, R. M. Alhebshi, A. G. Almakky, et al., Integrating artificial intelligence for improved security of IoT-drones through cyber-physical attack detection, *Frontiers in Computer Science*, Vol. 7, 2025, 1545282.

- [2]. A. Al-Sabbagh, A. El-Bokhary, S. El-Koussa, A. Jaber, et al., Enhancing UAV security against GPS spoofing attacks through a genetic algorithm-driven deep learning framework, *Information*, Vol. 16, Issue 2, 2025, 115.
- [3]. L. Alhoraibi, D. Alghazzawi, R. Alhebshi, Detection of GPS spoofing attacks in UAVs based on adversarial machine learning model, *Sensors*, Vol. 24, Issue 18, 2024, 6156.
- [4]. A. Almadhor, J. Baili, S. Alsubai, A. Al Hejaili, et al., CTDNN-Spoof: compact tiny deep learning architecture for detection and multi-label classification of GPS spoofing attacks in small UAVs, *Scientific Reports*, Vol. 15, Issue 1, 2025, 6656.
- [5]. T. Hickling, N. Aouf, P. Spencer, Robust adversarial attacks detection based on explainable deep reinforcement learning for UAV guidance and planning, *IEEE Transactions on Intelligent Vehicles*, Vol. 8, Issue 10, 2023, pp. 4381-4394.
- [6]. J. Zhou, M. Hu, C. Zhou, Z. Liu, et al., Research on GNSS spoofing detection and autonomous positioning technology for drones, *Electronics*, Vol. 14, Issue 15, 2025, 3316.
- [7]. L. Al-Soufi, T. Demirsoy, E. Gelal Soyak, An implementation-based study of the detection of and recovery from GPS spoofing attacks for unmanned aerial vehicles, *International Journal of Advances in Engineering and Pure Sciences*, Vol. 36, Issue 3, 2024, pp. 211-223.
- [8]. S. E. Meheretu, E. Nigussie, G. B. Gebremeskel, S. Y. Hailesilassie, A systematic literature review on spoofing and jamming approaches in unmanned aerial vehicles navigation, *Journal of Aerospace Technology and Management*, Vol. 17, 2025, e3425.
- [9]. D. Alsadie, J. Tian, Cybersecurity and artificial intelligence in unmanned aerial vehicles: Emerging challenges and advanced countermeasures, *IET Information Security*, Vol. 2025, 2025, 2046868.
- [10]. O. Y. Al-Jarrah, A. S. Shatnawi, M. M. Shurman, O. A. Ramadan, et al., Exploring deep learning-based visual localization techniques for UAVs in GPS-denied environments, *IEEE Access*, Vol. 12, 2024, pp. 113049-113071.
- [11]. L. A. Fagundes-Junior, K. B. de Carvalho, R. S. Ferreira, A. S. Brandão, Machine learning for unmanned aerial vehicles navigation: An overview, *SN Computer Science*, Vol. 5, Issue 2, 2024, 256.
- [12]. M. Adam, M. Hammoudeh, R. Alrawashdeh, B. Alsulaimy, A survey on security, privacy, trust, and architectural challenges in IoT systems, *IEEE Access*, Vol. 12, 2024, pp. 57128-57149.
- [13]. S. K. Khan, Cybersecurity modelling for the adoption and deployment of connected and automated vehicles, PhD Thesis, *University of Wollongong*, Wollongong, 2023.
- [14]. M. S. Islam, A. S. Mahmoud, T. R. Sheltami, AI-enhanced intrusion detection for UAV systems: A taxonomy and comparative review, *Drones*, Vol. 9, Issue 10, 2025, 682.
- [15]. J. P. A. Yaacoub, H. N. Noura, O. Salman, K. Chahine, Toward secure smart grid systems: Risks, threats, challenges, and future directions, *Future Internet*, Vol. 17, Issue 7, 2025, 293.
- [16]. A. Oun, K. Wince, X. Cheng, The role of artificial intelligence in boosting cybersecurity and trusted embedded systems performance: A systematic review on current and future trends, *IEEE Access*, Vol. 13, 2025, pp. 55258-55276.
- [17]. S. Silva, Secure edge AI for autonomous decision-making in unmanned aerial systems, in *Proceedings of the 2nd International Conference on Drones and Unmanned Systems (DAUS'26)*, 2026, pp. 189-190.

Advances in Signal Processing and Artificial Intelligence

Proceedings of the 2nd ASPAI' 2020 Conference



The proceedings contains all accepted and presented papers of both: oral and poster presentations at ASPAI' 2020 conference of authors from 23 countries. The coverage includes artificial neural networks, emerging trends in machine and deep learnings, knowledge-based soft measuring systems, artificial intelligence, signal, video and image processing.

Formats: hardcover (print book) and PDF (e-book), 264 pages
 ISBN: 978-84-09-21931-5, e-ISBN: 978-84-09-21930-8
 IFSA Publishing, 2020



https://www.sensorsportal.com/HTML/BOOKSTORE/ASPAI_2020_Proceedings.htm



Published by International Frequency Sensor Association (IFSA) Publishing, S. L., 2026
 (<https://www.sensorsportal.com>).

Civil-Military Knowledge Flows and the Diffusion of UAV Technologies under Economic Security Conditions

^{1,*} K. MIWA and ² K. SUMIKURA

¹ Science Technology and Innovation Policy Program, National Graduate Institute for Policy Studies (GRIPS), 7-22-1 Roppongi, Minato-ku, Tokyo 106-8677, Japan

² National Graduate Institute for Policy Studies (GRIPS), 7-22-1 Roppongi, Minato-ku, Tokyo 106-8677, Japan

¹ Tel.: + 81-3-6439-6000

* E-mail: doc24055@grips.ac.jp

Received: 28 Dec. 2025 /Revised: 8 April 2026 /Accepted: 15 April 2026 /Published: 28 April 2026

Abstract: Unmanned Aerial Vehicles (UAVs) have rapidly developed as emblematic dual-use technologies in which civilian and military sectors interact and jointly shape innovation trajectories. Recent intensifications of economic security policies and geopolitical tensions suggest that both knowledge flows and international technology diffusion in the UAV domain are being reconfigured. Building on our previous analysis presented at Defense and Aerospace Unmanned Systems 2026 (DAUS 2026), this study examines the relationship between national research structures and the international diffusion of attack-type UAVs. In particular, it pays attention to heterogeneity within attack-type UAV diffusion, including differences between armed UAVs and one-way loitering munitions, by combining bibliometric analysis of approximately 180000 Scopus-indexed UAV-related publications (2001–2025) with the CNAS Drone Proliferation Dataset.

The analysis identifies a marked increase in Civil-to-Military (C2M) knowledge flows since the late 2010s, indicating the growing importance of civilian-led technological inputs in military UAV development. This shift corresponds with a structural transformation in global diffusion. Mean-shift analysis detects a first major breakpoint in 2012, when the combined supplier share of China and Turkey rose from near zero to a regime average of 0.52, followed by a second breakpoint in 2018, when the average annual number of adoption events increased from 2.7 to 24.8.

These findings suggest a staged transformation in which changes in supplier-side institutional and industrial structures preceded the subsequent expansion of international UAV adoption. Importantly, the results suggest a sequential transformation in which supply-side structural changes precede and condition subsequent diffusion expansion. Recent geopolitical shocks appear to have accelerated this process by shifting diffusion from unarmed baseline systems toward attack-capable platforms, including the rapid expansion of non-traditional suppliers such as Iran in the loitering munition segment. Overall, the study identifies a systematic alignment between (non-causal) changes in national research structures and patterns of technology diffusion, highlighting the growing role of military organizations as absorptive users of civilian-led innovation under contemporary economic security conditions.

Keywords: UAV diffusion, Dual-use technology, Economic security, Bibliometric analysis, Civil–military knowledge flow.

1. Introduction

Unmanned Aerial Vehicles (UAVs) represent one of the most rapidly diffusing military technologies [1], with the proliferation of reusable armed UAVs and one-way loitering munitions fundamentally transforming modern military operations [2]. This

diffusion process represents a staged transformation often catalyzed by geopolitical shocks. Simultaneously, UAV technologies are becoming pervasive in civilian domains such as agriculture, logistics [3], and disaster response [4, 5], further blurring the technological boundaries between military and civilian sectors.

As geopolitical competition intensifies, UAV-related technologies have increasingly become subject to economic security policies, including export controls, research-security regimes, and procurement restrictions [6]. This perspective aligns with the concept of “weaponized interdependence,” in which states leverage their position within global technological and economic networks to exert strategic influence through regulatory and access-based mechanisms [6]. Technological development in such fields does not proceed in a linear manner from basic research to application, but evolves through iterative interactions among knowledge production, use environments, and institutional conditions [7].

In the development of such dual-use technologies, the direction and intensity of knowledge flows – specifically transfers from civilian research to military applications and vice versa – play a decisive role in shaping technological trajectories [6, 8]. However, systematic empirical analysis focusing explicitly on these civil–military knowledge flows within the UAV domain remains limited, particularly regarding the military domain's role as an absorptive user of civilian-led innovations during the shift from unarmed to attack-type platforms.

Building on our previous conference paper presented at DAUS 2026 [9], the present study extends the earlier analysis in three respects. First, it disaggregates attack-type UAV diffusion into reusable armed UAVs and one-way loitering munitions, allowing a closer examination of heterogeneity in supplier structure and diffusion pathways. Second, it explicitly distinguishes Civil-to-Military (C2M) and Military-to-Civil (M2C) citation flows in order to

analyze the directional structure of knowledge circulation. Third, it more systematically integrates bibliometric evidence with the CNAS diffusion dataset, linking changes in citation structures to shifts in supplier composition and proliferation volume.

To do so, the study combines two types of evidence. First, it analyzes directional citation flows – Civil-to-Military (C2M) and Military-to-Civil (M2C) – within a large corpus of Scopus-indexed UAV-related publications (2001–2025). Second, it relates these knowledge-flow patterns to updated international diffusion data for attack-type UAVs (1995–2023) based on the CNAS Drone Proliferation Dataset [10]. This design allows us to examine whether changes in knowledge circulation correspond to structural shifts in supplier composition and proliferation volume.

2. Data and Methodology

The analysis utilizes a bibliometric dataset of 182291 UAV-related publications indexed in Scopus over the period from 2001 to 2025, based on records retrieved on 1 December 2025. Fig. 1 provides an overview of the Scopus search and rule-based classification procedure, while the full search query and keyword dictionaries are reported in Appendix A for transparency and reproducibility. To improve transparency, the rule-based search logic and classification criteria used for military, civilian, and dual-use categories are described more explicitly in this section, and the full search query and keyword dictionaries are provided in Appendix A.

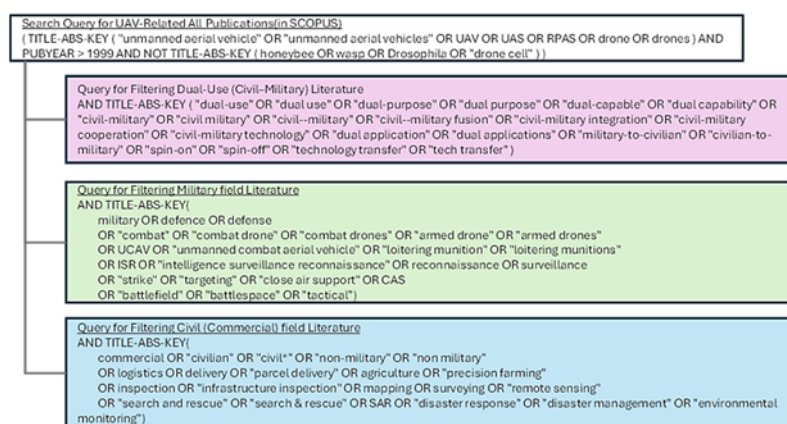


Fig. 1. Overview of the Scopus search and rule-based classification procedure.

The classification is based on keyword filters applied to titles, abstracts, and author keywords, and is intended as an analytical coding procedure for identifying directional knowledge relationships rather than as a definitive classification of all UAV-related publications. It is not designed to exhaustively label the full literature, but to provide directional signals for citation-based analysis, which serves as the primary

basis of the analysis. In the present study, the dual-use category refers to records explicitly identified through dual-use-related terms, while directional citation analysis is used separately to capture cross-domain knowledge relationships. Using this procedure, we classified the literature into military (23024), civilian (57720), and dual-use (332) categories, building on the classification approach developed in our

earlier study [9]. The remaining approximately 100000 publications primarily consist of foundational research that lacks domain-specific labels but serves as a technological resource for both sectors. Although only 332 papers are explicitly labeled as “dual-use,” the much larger volume of C2M citations suggests that the dual-use character of UAV technology is more visible in cross-domain knowledge absorption than in formal publication labels. Two directional citation indicators were then constructed: C2M citations, representing military absorption of civilian research (“spin-on”), and M2C citations, representing civilian spillovers from military research (“spin-off”).

In parallel, we utilized the CNAS Drone Proliferation Dataset, which tracks drone transfers from 1995 to 1 September 2023 [10], to analyze the international spread of attack-type UAV systems across countries and over time. A small number of earlier entries (1992–1993) appear in the raw dataset; however, these correspond to non-armed UAV observations and fall outside the attack-type UAV subset used in this study, and therefore do not affect the results reported here. In this study, following the CNAS codebook definitions, the category of “attack-type UAVs” is disaggregated into two distinct platform types: reusable armed UAVs (weapons-capable platforms) and one-way loitering munitions (kamikaze drones). To ensure transparency in platform selection, Appendix Table A1 lists the representative series included in the diffusion analysis, grouped by major supplier country. The appendix table is intended to clarify the operational scope of the dataset rather than to provide an exhaustive inventory of all variants and sub-models.

The dataset includes representative platforms that can be broadly grouped into two distinct diffusion models. The first is a high-end, state-centric model led by the United States and Israel, centered on MALE-type and other high-performance UAVs such as the MQ-1 Predator, MQ-9 Reaper, Heron TP, and Hermes 450. These systems typically rely on state-led R&D and procurement, long-endurance operations, beyond-line-of-sight control via satellite communications, advanced ISR payloads including EO/IR and SAR, and, in armed variants, precision-strike weapons such as the AGM-114 Hellfire. Within this broader diffusion landscape, loitering munitions such as the Harpy, Harop, SkyStriker, and the HERO series represent a more adaptive and cost-conscious segment. These systems generally emphasize loitering capability, flexible target engagement, and lower-cost precision attack, often using combinations of anti-radiation seekers or EO/IR-based man-in-the-loop guidance.

The second is an export-oriented and adaptive model, represented by emerging systems from China, Turkey, and Iran. These suppliers emphasize accessibility, exportability, and rapid iteration, including China’s CH and Wing Loong series, Turkey’s Bayraktar TB2, Akinci, and ANKA series, and Iran’s Ababil and Shahed series, which have

recently established significant footprints outside traditional Western networks.

By combining these bibliometric indicators with documented international platform transfers, this study empirically connects shifts in national research structures with the staged transformation of global technological proliferation. This methodology allows for a systematic examination of how critical structural breaks – such as the 2012 supplier share shift and the 2018 acceleration in adoption – correlate with the military domain’s evolving capacity to absorb civilian-led innovations.

3. Results

3.1. International Diffusion of Attack-Type UAVs

The global diffusion of attack-capable UAVs has undergone a structural shift described as a staged transformation, where the widespread adoption of unarmed systems provided the foundation for a subsequent surge of attack-type platforms (see Fig. 2). This process was significantly catalyzed by geopolitical shocks, with a structural breakpoint in 2018, when the annual average number of adoption events increased from 2.7 to 24.8 (see Fig. 7b for details). While the 2000s were dominated by the U.S.–Israel-led high-end/state-centric model, the late 2010s have been defined by China- and Turkey-centered export-oriented/adaptive models. This shift is characterized by greater accessibility and a rapid increase in the volume of international transfers, where competitive advantage is increasingly associated with development speed and the capability to incorporate operational feedback.

To clarify how this overall shift was distributed across suppliers and platform types, the next section disaggregates diffusion patterns by country and system category.

3.2. Regional Diffusion Patterns and Platform Dominance

Taken together, Figs. 3-5 show diffusion at three connected levels: supplier composition, regional network expansion, and platform-level transfer volume. Fig. 3 disaggregates global UAV proliferation by supplier country and platform type. The figure shows that unarmed UAV diffusion was initially led by the United States and Israel, whereas the later spread of armed UAVs increasingly involved China and Turkey. In the loitering munition segment, Israel appears as an early supplier, while Iran emerged rapidly in the late 2010s and early 2020s. These country-level differences indicate that the diffusion of UAV technologies has become more diversified over time, with distinct supplier profiles depending on platform type.

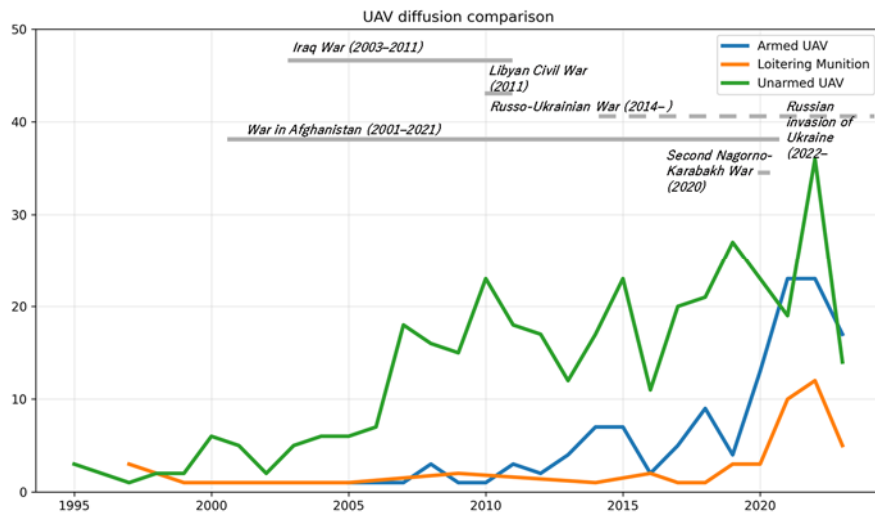


Fig. 2. Comparison of UAV diffusion by type (1995–2023).

Note: The figure illustrates the temporal evolution of adoption events across unarmed UAVs, armed UAVs, and loitering munitions. The results indicate a gradual transition from early commercial UAV diffusion to increased adoption of armed UAVs and loitering munitions in recent years.




Fig. 3. Global Proliferation of UAVs by supplier country.

Note: To minimize gaps in the time-series data, the annual number of UAV adoption events prioritizes the "Year of Order" from the CNAS dataset. Where this data is missing, it is supplemented by the "Year of First Delivery," followed by the "Year Identified, Completed, or Canceled" and the "Year of Interest" to ensure maximum data coverage.

The expanded regional supplier–seeker network reveals distinct geopolitical footprints among the six major suppliers (see Fig. 4). While the United States and Israel maintain strong ties with seekers in Europe and the Asia-Pacific, China and Turkey have established particularly strong positions in the Middle

East and Africa through their export-oriented/adaptive models. Notably, Iran has developed distinctive supply routes to the Americas (specifically Venezuela) and East Africa (e.g., Ethiopia and Sudan), suggesting a diversification of supply channels outside traditional Western-centered networks.




Fig. 4. Expanded Regionalized Supplier-Seeker Network of Attack-type UAV Diffusion.

Transfer volume data by series further clarifies the drivers of this diffusion (see Fig. 5). The Bayraktar TB2 (Turkey) stands out as a particularly notable case in the adaptive category for reusable armed UAVs, with more than 40 documented transfer events, second only to Israel’s long-standing Hermes/Heron series.

Furthermore, specifically in the loitering munition segment, Iranian series such as the Ababil and Shahed series exhibit transfer volumes comparable to established Chinese platforms such as the Wing Loong series and have become primary drivers of proliferation in recent conflict zones.




Fig. 5. Representative Attack-type UAV Series by Supplier.

These observations suggest that the staged transformation identified in this study is driven not only by increased transfer volume, but also by the emergence of highly competitive non-traditional supplier models that rapidly weaponize the civilian-led technological base while prioritizing accessibility and regional strategic alignment.

3.3. Military/Civil Publication Ratio (Fig. 6)

To interpret the knowledge conditions underlying this diffusion shift, we next examine whether the

research base itself has become more civilian-led over time.

Globally, UAV-related publications have increased steadily since the early 2000s, with particularly rapid growth after 2014. Analysis of the Military/Civil publication ratio reveals a global trend toward civilian-led knowledge production, as the ratio has declined in most countries since 2001. This highlights the growing importance of commercial technologies – such as sensors, autonomy, and AI – as foundational inputs into UAV development. This suggests that the knowledge base has structurally shifted toward civilian-dominated production, increasing the potential for military absorption.




Fig. 6. Military/Civil UAV Publications Ratio by Country (5-year periods).

This civilian-led knowledge base provides a reservoir that military domains can absorb during periods of strategic necessity. Ukraine is an important exception, showing a more security-driven knowledge structure from the mid-2000s onward.

3.4. Structural Break Analysis of Proliferation Dynamics (Fig. 7)

To identify structural changes in UAV proliferation, we applied a two-regime Poisson mean-shift model to annual adoption data and a mean-shift model to supplier shares for the period 2000 to 2023. The estimation procedure follows a likelihood-based segmentation approach, in which candidate breakpoints are evaluated by comparing model fit across alternative partitions of the time series using a penalized criterion (BIC). These procedures are used here as exploratory tools for detecting regime-level discontinuities in the data. The analysis identifies two critical breakpoints. First, a shift in the supplier base occurred in 2012, after which the combined average share of China and Turkey for the subsequent regime rose from near zero to 0.52 (52 %) (Fig. 7a). Second, a major break in proliferation volume was detected in 2018, when the average annual number of proliferation events shifted from 2.7 to 24.8 (Fig. 7b). The identified breakpoints are interpreted as empirically observed structural shifts without causal identification rather than as definitive causal estimates.

Together, these results demonstrate the staged nature of this transformation in UAV diffusion, in which structural changes in the supplier ecosystem around 2012 established the supply-side capacity that was subsequently activated by geopolitical shocks. These shocks acted as catalysts, triggering the sharp acceleration of adoption after 2018, as seekers shifted

from unarmed systems to battle-proven armed UAVs and loitering munitions. This transition underscores the military domain's evolving role as an absorptive user capable of rapidly weaponizing the civilian technological base in response to regional conflicts.

3.5. Citation Relationships and Knowledge Flows (Fig. 8)

To assess whether this diffusion shift was accompanied by a corresponding reorientation in knowledge circulation, we finally examine directional citation flows.

Citation analysis reveals a clear asymmetry in knowledge circulation across national innovation ecosystems. Globally, C2M citation flows have increased significantly (reaching 91642 citations), indicating that the civilian domain has become a primary source of technological inputs for military UAV development.

In contrast, M2C citation flows exhibit more stagnant and heterogeneous patterns across major UAV-producing countries.

These observations suggest that while knowledge absorption from civilian technology (spin-on) is accelerating, it creates a pervasive technological reservoir that enables the staged transformation of the ecosystem, allowing the military domain to act as an "absorptive user" during strategic shocks. Conversely, spillovers back into the civilian domain (spin-off) remain uneven and dependent on specific institutional configurations.

These citation patterns should, however, be interpreted with caution, since military knowledge production is not always fully visible in open publications [11] and may therefore be underrepresented in bibliometric data.






Fig. 7. Structural breaks analysis of proliferation events and armed UAV supplier shares (2000–2023).



(a) M2C citation flows



(b) C2M citation flows

Fig. 8. M2C and C2M citation flows in UAV-related research (2001–2025).

4. Discussion

The main implication of these findings is that the strategic value of UAV-related innovation depends not only on indigenous technological performance, but also on how effectively civilian-led knowledge can be absorbed and mobilized under conditions of geopolitical stress. This dynamic is particularly pronounced in UAV technologies, where modular architectures, rapid iteration cycles, and operational

feedback loops enable fast translation of civilian components into deployable military systems.

Recent empirical evidence from the 2018 breakpoint suggests that this diffusion is not merely linear [7], but follows a staged transformation path where geopolitical shocks act as critical catalysts. These shocks bridge the gap between the civilian-led knowledge base and the rapid proliferation of attack-type platforms, including both reusable armed UAVs and loitering munitions. This is further

evidenced by the emergence of non-traditional suppliers like Iran, which has successfully weaponized established technology bases to capture significant shares in the loitering munition segment.

These results further suggest that under economic security constraints, the capacity to iterate and absorb external knowledge may become as decisive as the ability to translate indigenous technological performance into operational capabilities [12, 13]. This underscores the importance of adaptive supplier models, like those of Turkey and China, which established the necessary supply-side capacity (achieving a 0.52 market share) following the 2012 structural shift. Ultimately, an ecosystem's strategic impact is determined by its primary transmission channels, which dictate the speed at which accumulated civilian-led innovations are absorbed and converted into military capabilities during periods of strategic necessity.

These dynamics are consistent with broader empirical evidence that military-related innovations can diffuse as effectively as civilian ones, particularly when they function as general-purpose technologies that overcome institutional barriers [11]. Furthermore, the development of such critical technologies has been shown to generate substantial spillover effects, stimulating R&D activities across the wider innovation ecosystem, even among organizations not directly involved in military projects [14].

For Japan, this implies that economic security policy [15] should address not only the protection of critical technologies, but also the institutional constraints on civil–military integration, which may limit the effective absorption of rapidly evolving civilian UAV technologies and potentially create a structural disadvantage in adaptive capability.

5. Conclusion

This study yields three main findings. First, UAV research has increasingly shifted toward a structure in which military development draws on a growing civilian-led knowledge base. Second, the international diffusion of attack-type UAVs accelerated sharply from the late 2010s, following an earlier structural shift in the supplier base. Third, these results suggest a systematic correspondence between changes in knowledge circulation and patterns of technology diffusion under economic security conditions.

This study is also subject to the limitation that military knowledge is less likely to be fully disclosed in open publications, which may lead to an underestimation of military knowledge production in bibliometric and citation-based analysis. The observed relationship between knowledge circulation and technology diffusion should therefore be interpreted as a structural association rather than a causal effect.

Taken together, these findings highlight the importance of understanding knowledge circulation when analyzing the evolution of dual-use technologies under economic security conditions.

References

- [1]. D. C. Mowery, Military R&D and innovation, in Handbook of the Economics of Innovation (B. H. Hall, N. Rosenberg, Eds.), Vol. 2, Elsevier, 2010, pp. 1219-1256.
- [2]. V. W. Ruttan, Is War Necessary for Economic Growth? Military Procurement and Technology Development, Oxford University Press, 2006.
- [3]. OECD Science, Technology and Innovation Outlook 2016, OECD Publishing, Paris, 2016.
- [4]. OECD Science, Technology and Innovation Outlook 2018: Adapting to Technological and Societal Disruption, OECD Publishing, Paris, 2018.
- [5]. OECD Science, Technology and Innovation Outlook 2021: Times of Crisis and Opportunity, OECD Publishing, Paris, 2021.
- [6]. H. Farrell, A. L. Newman, Weaponized interdependence: How global economic networks shape state coercion, International Security, Vol. 44, Issue 1, 2019, pp. 42-79.
- [7]. N. Rosenberg, Critical issues in science policy research, Science and Public Policy, Vol. 18, Issue 6, 1991, pp. 335-346.
- [8]. R. Adner, R. Kapoor, Value creation in innovation ecosystems: How the structure of technological interdependence affects firm performance in new technology generations, Strategic Management Journal, Vol. 31, Issue 3, 2010, pp. 306-333.
- [9]. K. Miwa, K. Sumikura, Policy and institutional drivers of UAV innovation ecosystems under economic security constraints, in Proceedings of the 2nd International Conference on Drones and Unmanned Systems (DAUS'26), 2026, pp. 59-65.
- [10]. M. Campbell, Drone proliferation dataset, Center for a New American Security (CNAS), 2024.
- [11]. J. Schmid, The diffusion of military technology, Defence and Peace Economics, Vol. 29, Issue 6, 2018, pp. 595-613.
- [12]. D. J. Teece, Explicating dynamic capabilities: The nature and microfoundations of (sustainable) enterprise performance, Strategic Management Journal, Vol. 28, Issue 13, 2007, pp. 1319-1350.
- [13]. A. Gilli, M. Gilli, Why China has not caught up yet: Military-technological superiority, systems integration, and the challenges of imitation, reverse engineering, and cyber-espionage, International Security, Vol. 43, Issue 3, 2019, pp. 141-189.
- [14]. A. Kotani, K. Nakajima, T. Okazaki, Y. U. Saito, The impact of military technology development on innovation: Evidence from pre-war and wartime Japanese secret patents, RIETI Discussion Paper Series No. 25-E-093, 2025.
- [15]. Ministry of Foreign Affairs of Japan, The three principles on transfer of defense equipment and technology, 2023, https://www.mofa.go.jp/fp/nsp/page1we_000083.html

Appendix

Table A1 summarizes the representative attack-type UAV series included in the network and diffusion analysis. The table groups platforms by major supplier countries and distinguishes between armed UAVs and loitering munitions where relevant. Its purpose is to clarify the analytical coverage of the

study by showing which platform families were treated as representative cases in the construction of the dataset. The table should therefore be read as an operational classification for analysis, rather than as a complete catalog of all exported UAV variants.

UAV-related publications into military, civilian/commercial, and dual-use categories. The purpose of this appendix is to improve transparency and reproducibility by presenting the search logic and classification procedure in text form.

Appendix A. Full Scopus Search Query and Operational Classification Rules

This appendix provides the full Scopus search query and the operational rules used to classify

A1. Search Field and Retrieval Date

The bibliometric corpus was retrieved from Scopus on **1 December 2025**. The search was applied to **titles, abstracts, and author keywords (TITLE-ABS-KEY)**.

Table A1. Representative attack-type UAV series included in the diffusion analysis (CNAS data).

United States	Israel	China	Turkey	Iran
Armed UAV	Armed UAV	Armed UAV	Armed UAV	Loitering Munition
• MQ-1 Predator	• Hermes 450	• CH-3	• Bayraktar TB2	• Ababil-2T
• MQ-9 Reaper	• Heron TP	• CH-4	• Akinci	• Ababil-T
		• CH-4B	• ANKA	• Jenin
Loitering Munition	Loitering Munition	• CH-5	• ANKA-S	• Khatif-1
• Phoenix Ghost	• Harop	• CH-92	• Aksungur	• Khatif-2
• Switchblade 300	• Harpy	• CH-95	• Karayel	• Murad-5
• Switchblade 600	• Skystriker	• Wing Loong I	• Songar	• Samad-2
	• Hero-30	• Wing Loong II		• Samad-3
	• Hero-120	• Wing Loong ID	Loitering Munition	• Shahed
	• Mini Harpy	• CR500 Golden Eagle	• Kargu	• Shahed-131
	• Orbiter 3	• WZ-7 Xianglong		• Shahed-136
	• Orbiter-1K	• Unnamed System		• Shihab
	• SpyX	(Unknown name)		

Note: The table reports representative platform families used for the diffusion analysis. Minor variants and closely related sub-models were aggregated where appropriate to maintain analytical comparability and visual readability.

A2. Broad UAV Search Query

The broad UAV corpus was constructed using the following Scopus query: [TITLE-ABS-KEY ("unmanned aerial vehicle" OR "unmanned aerial vehicles" OR UAV OR UAS OR RPAS OR drone OR drones) AND PUBYEAR > 1999 AND NOT TITLE-ABS-KEY (honeybee OR wasp OR Drosophila OR "drone cell")].

This query was designed to capture the broad UAV-related literature across both military and civilian/commercial contexts.

A3. Domain-Specific Query Sets

After retrieval of the broad UAV corpus, rule-based keyword filters were used to identify domain-specific subsets.

A3.1. Dual-Use Query Set

Records were classified as dual-use when they were explicitly identified through the following dual-use-related query set: [AND TITLE-ABS-KEY

("dual-use" OR "dual use" OR "dual-purpose" OR "dual purpose" OR "dual-capable" OR "dual capability" OR "civil-military" OR "civil military" OR "civil--military" OR "civil--military fusion" OR "civil-military integration" OR "civil-military cooperation" OR "civil-military technology" OR "dual application" OR "dual applications" OR "military-to-civilian" OR "civilian-to-military" OR "spin-on" OR "spin-off" OR "technology transfer" OR "tech transfer")].

In the present study, the dual-use category refers to records explicitly identified through dual-use-related terms. Cross-domain knowledge relationships were not inferred from keyword overlap alone, but were examined separately through directional citation analysis.

A3.2. Military-Oriented Query Set

Records were classified as military-oriented when they matched the following military-related query set: [AND TITLE-ABS-KEY(military OR defence OR defense OR "combat" OR "combat drone" OR "combat drones" OR "armed drone" OR "armed drones" OR UCAV OR "unmanned combat aerial

vehicle" OR "loitering munition" OR "loitering munitions" OR ISR OR "intelligence surveillance reconnaissance" OR reconnaissance OR surveillance OR "strike" OR "targeting" OR "close air support" OR CAS OR "battlefield" OR "battlespace" OR "tactical"]].

This query set was intended to identify publications associated with defense, combat, reconnaissance, targeting, strike, surveillance, and related military mission contexts.

A3.3. Civilian/Commercial-Oriented Query Set

Records were classified as civilian/commercial-oriented when they matched the following non-military query set: [AND TITLE-ABS-KEY(commercial OR "civilian" OR "civil*" OR "non-military" OR "non military" OR logistics OR delivery OR "parcel delivery" OR agriculture OR "precision farming" OR inspection OR "infrastructure inspection" OR mapping OR surveying OR "remote sensing" OR "search and rescue" OR "search & rescue" OR SAR OR "disaster response" OR "disaster management" OR "environmental monitoring")].

This query set was intended to identify publications associated with logistics, agriculture, mapping, inspection, monitoring, disaster response, infrastructure, and related civilian or commercial applications.

A4. Operational Classification Procedure

The classification was implemented according to the following procedure:

1. All records were first required to be included in the broad UAV corpus;
2. Records matching the military-oriented query set were classified as military-oriented;
3. Records matching the civilian/commercial-oriented query set were classified as civilian/commercial-oriented;
4. Records matching the dual-use query set were classified as dual-use;
5. Records not matching these domain-specific query sets were retained in the broad UAV

corpus as foundational or enabling research, but were not assigned to the three focal categories.

This procedure was designed as an **analytical coding framework** for identifying directional knowledge relationships across domains, rather than as a definitive substantive labeling of all UAV-related publications.

A5. Interpretation of the Dual-Use Category

The dual-use category should be interpreted cautiously. The number of publications identified through explicit dual-use-related terms does not represent the full extent of dual-use knowledge in the UAV field. Rather, it indicates the number of records formally labeled in such terms. In this study, broader cross-domain knowledge relationships are examined through directional citation analysis, especially Civil-to-Military (C2M) and Military-to-Civil (M2C) citation flows.

A6. Interpretation of Unclassified Publications

A substantial number of UAV-related publications were not assigned to the military, civilian/commercial, or dual-use categories. These records primarily consist of foundational or enabling research lacking explicit domain-specific labels. They were retained in the broad UAV corpus because they still form part of the wider technological knowledge base relevant to both civilian and military development.

A7. Limitation of the Classification Procedure

A key limitation of this bibliometric approach is that military research is less likely to be fully disclosed in open publications than civilian research. As a result, military-oriented knowledge production may be underrepresented in the observable literature. The present classification should therefore be interpreted with this structural visibility bias in mind.



8th International Conference on Advances in Signal Processing and Artificial Intelligence (ASPai' 2026)

ASPai
MALLORCA • 2026

24-26 June 2026 • Palma de Mallorca (Balearic Islands), Spain

CALL FOR PAPERS

DEADLINES

Submission (2-page abstract):
still open

Notification of acceptance:
in 2 weeks after submission

Registration:
as soon as possible

Camera ready (4-6 pages
paper): **30 May 2026**

SPECIAL SESSIONS

Participants are welcome to organize and manage special sessions during the conference. Each session will contain 4-6 papers in a related field as specified above.

Session organizers will get:

- Certificate of Appreciation
- Free Registration for the Event
- Special Publishing Theme within Conference Proceedings

Advances in AI and signal processing are driving market growth, enabling better technologies for drones, robotics, and autonomous cars. Sensor manufacturers increasingly use machine learning for data analysis, which has become easier as hardware shrinks and sensors become cheaper. This expands IoT applications but creates challenges in managing the growing data volumes. Traditional processing methods can no longer handle the data efficiently.

The ASPai conferences aim to address these issues by offering a platform for discussing AI and signal processing technologies. The 2025 event, organized by IFSA, will coincide with the International Conference on AI in Medicine and Healthcare (AiMH' 2026). It will feature talks, presentations, and networking events, creating opportunities for collaboration among researchers and industry professionals.

ARTIFICIAL INTELLIGENCE

- AI Algorithms
- Intelligent System Architectures
- Hybrid Intelligent Systems
- Artificial Neural Networks
- Expert Systems
- Intelligent Networks
- Parallel Processing
- Pattern Recognition
- Pervasive Computing and Ambient Intelligence
- Programming Languages Artificial Intelligence
- Soft Computing and Applications
- Artificial Intelligence Tools & Applications
- CAD Design & Testing
- Computer Vision and Speech Recognition
- Fuzzy Logic, Systems and learning
- Computational Theories, Learning and Intelligence
- Computational Neuroscience
- Soft Computing Theory and Applications
- Software & Hardware Architectures
- Web Intelligence Applications & Search
- Ambient Intelligence
- Artificial Immune Systems
- Autonomous and Ubiquitous Computing
- Bayesian Models and networks
- Data Fusion
- Distributed AI
- Machine Learning
- Deep Learning
- Learning and Adaptive Sensor Fusion
- Multisensor Data Fusion Based on Neural and Fuzzy Techniques
- Applied Artificial Intelligence
- Self-Organising Networks
- Virtual and Augmented Reality
- Emotion Detection

SIGNAL PROCESSING

- Signal Processing Theory and Methods
- Sensor Array and Multichannel Signal Processing
- Multivariable Sensor Systems
- Digital Signal Processing
- Image and Video Processing
- Multichannel Signal Processing
- Biomedical Signal Processing
- Biometrics
- Chaotic and Fractal Systems
- Computer Vision and Pattern Recognition
- Image and Video Processing
- Information Theory and Coding
- Speech Signal Processing
- Signal Processing for Communications and Networking
- Radar and Sonar Signal Processing
- Satellite Signal Processing
- Nonlinear Signal Processing
- Statistical Signal Processing



aspai@sensorsportal.com

<https://aspai-conference.com/>





COMMITTEE

Chairman:

Prof., Dr. Sergey Y. Yurish (IFSA, Spain)
e-mail: aspai@sensorsportal.com

Advisory Chairman:

Prof., Dr. Hojjat Adeli (The Ohio State University, USA)

Organizing Committee:

Dr. Mobyen Uddin Ahmed (Mälardalen University, Sweden)
Dr. Sergey Grosman (Siemens PPAL, Germany)
Prof. Sandeep Singh Sengar (Cardiff Metropolitan University, UK)

Conference Manager:

Mrs. Tetyana Zakharchenko (IFSA Publishing, S.L., Spain)

VENUE

The ASPAI' 2026 will take place in the Melia Palma Marina Hotel, Palma de Mallorca (Balearic Islands), Spain.




The Hotel Melia Palma Marina is located on the seafront promenade, just 500 m from Palma de Mallorca's 14th-century Bellver Castle. Palma Cathedral is only 1.5 km away. This grand hotel features a spa centre, a seasonal outdoor swimming pool and sun terrace with views of the sea and the Bay of Palma. Each air-conditioned room has a balcony and free Wi-Fi. The hotel is situated in 9 km from Palma de Mallorca Airport.

Address: Av. de Gabriel Roca, 29, 07014 Palma, Illes Balears, Spain



ONE EVENT - THREE DIFFERENT PUBLICATIONS !

1. All registered abstracts will be published in the conference proceedings with the ISBN, ISSN and DOI.
2. Authors will be invited to submit full-page extended papers to one of special journals' issues:
 - Sensors & Transducers (ISSN 2306-8515, e-ISSN 1726-5479) by IFSA Publishing
 - Integrated Computer-Aided Engineering (ISSN: 1069-2509, e-ISSN: 1875-8835) by IOS Press
 - International Journal of Neural Systems (ISSN: 0129-0657, e-ISSN: 1793-6462) by World Scientific
3. Authors will be also invited to submit their book chapters to open access books' chapters for the Book Series on 'Advances in Signal Processing', Vol. 4, 'Advances in Artificial Intelligence', Vol. 3 or 'Advances in Intelligent Systems', Vol. 4, which will be published by IFSA Publishing at the end of 2026.



Call for Chapters

Your Chapter may be in the next Volume of the popular Book Series

ADVANCES IN SENSORS

9

Open Access Book Series

Since 2012



- Physical, chemical and biosensors
- Sensor uncertainty, accuracy, calibration
- Smart sensors and systems
- Intelligent sensors and systems
- Artificial Intelligence enabled sensors
- Virtual and soft sensors
- Sensor instrumentation
- Sensor buses and networks
- Signal processing, sensor fusion
- Nanosensors
- Photonic sensors
- Sensor interfacing and signal conditioning
- Materials for sensors
- MEMS, BioMEMS and NEMS
- Remote Sensors and Telemetry
- Sensor packages and reliability

https://sensorsportal.com/ifsa_publishing.html

