

TDP: Black Bee Drones (Brazil), Indoors Category

Ana Clara Alves Guilherme¹, Gabriel Dias Caproni², Guilherme de Paula Pinheiro³, João Victor Moreira de Oliveira⁴, Luis Filipe Laurino Junqueira Pedras⁵ and Bruno Guazzelli Batista⁶.

¹Federal University of Itajubá, MG, Brazil (UNIFEI)

ABSTRACT

The development of autonomous UAVs for IMAV 2025 required the integration of robust mechanical, hardware, and software solutions to address the challenges of indoor missions, such as GPS-denied navigation, variable lighting conditions, and obstacle-rich environments. The drone was designed with a lightweight yet resistant carbon fiber structure, advanced embedded computing using the NVIDIA Jetson Orin Nano Super, and a modular set of sensors including the Intel RealSense D435i, IMX219 camera, Pixhawk 6C Mini, and TF-Luna LiDAR. On the software side, a GPU-accelerated Visual SLAM pipeline based on Isaac ROS was implemented for precise navigation and positioning, while deep learning-based detection modules (YOLOv11) were used for tunnel gate alignment and autonomous landing under adverse conditions. The integration of these subsystems provides a reliable, high-performance platform capable of fulfilling the competition's requirements and simulating real-world search-and-rescue scenarios.

1 INTRODUCTION

Unmanned Aerial Vehicles (UAVs) have become increasingly important for indoor applications such as inspection, monitoring, and search-and-rescue, where GPS is unavailable and environmental conditions impose additional challenges. The International Micro Air Vehicle (IMAV) 2025 competition serves as a testing ground for innovative solutions in this context, pushing teams to design systems that combine mechanical robustness, advanced onboard computing, and intelligent autonomy.

The Black Bee Drones team addressed these challenges by developing a drone capable of performing indoor missions, including tunnel navigation and precise landing in low-visibility conditions. To achieve this, the project was structured around three main aspects: a carbon fiber-based mechanical frame designed for lightweight strength and vibration mitigation; a hardware architecture centered on the NVIDIA Jetson Orin Nano Super for high-performance onboard processing and a software system divided into two main axes: Visual SLAM for navigation and positioning, and deep learning-based detection modules for environmental interaction.

This paper details the main design choices and solutions adopted in the mechanical, electrical, electronic, and software systems, highlighting their integration into a cohesive UAV platform prepared to address the specific requirements of IMAV 2025 missions.

2 MECHANICAL SYSTEM

2.1 Previous Issues

Since the foundation of Black Bee Drones, the team has stood out by developing and manufacturing its own frames, prioritizing innovative solutions tailored to the specific needs of each project. In past projects, the drone structures were built from a sandwich composite consisting of carbon fiber, PVC foam, and epoxy resin. The main advantage of this configuration was its low weight, with frames weighing as little as approximately 150 grams. However, during flight tests and mission simulations, the need to increase mechanical stiffness became evident. The arms showed considerable deflections under aerodynamic loads, directly compromising flight stability. In addition, the material proved to be impact-sensitive, with frequent fractures upon crashes, often irreparable, which required complete replacement of the structure. Thus, what initially seemed advantageous in terms of cost proved unfeasible due to low durability and the recurring need for reconstruction, driving the search for more resistant and efficient alternatives.

Moreover, the mounts were manufactured in PLA (Polylactic Acid), a material easy to process but with low impact resistance, little flexibility, poor fatigue resistance, and a glass transition temperature around 328,15–333,15 K. These limitations resulted in frequent fractures, deformations under heat, and low reliability under real operating conditions, requiring recurrent replacements.

2.2. Solutions

In response to these challenges, the team carried out a complete redesign of the drone's structure for IMAV 2025. The new project adopted carbon fiber tubes for the arms of the frame and the landing gear, due to the combination of low weight and high mechanical strength [1]. The upper and lower plates were manufactured from

sheets of the same material, ensuring structural rigidity and stability during flight, while reducing deformations caused by aerodynamic loads and minimizing impact during landings. Carbon fiber also contributes to vibration absorption, an essential factor to mitigate fatigue and ensure higher control accuracy [1]. To further optimize performance, strategic holes were added to the plates, reducing weight without compromising strength, while improving airflow for cooling internal devices and facilitating cable routing and component assembly.

The arm and camera mounts were produced using 3D printing in ABS (Acrylonitrile Butadiene Styrene), a polymer with a glass transition temperature (T_g) of approximately 378,15 K [2-3]. This characteristic is essential to withstand proximity to components that reach high temperatures, such as motors. Furthermore, ABS's good impact resistance and flexibility allow these mounts to act as an intermediate layer between components and the carbon fiber structure, absorbing part of the vibrations and reducing their transmission to the frame [4].

Finally, the landing gear supports were manufactured in TPU (Thermoplastic Polyurethane), a material selected for its high flexibility and impact absorption capability [5-6]. These properties allow the dissipation of energy from hard landings, minimizing the risk of damage to the structure. TPU also offers excellent resistance to deformation and bending fatigue, returning to its original shape after repetitive stresses, which ensures greater durability and protection for both the drone and the contact surface [7-8].

Table 1 presents a comparison of the previously discussed mechanical properties of PLA, ABS, and TPU.

Characteristics	PLA	ABS	TPU
Flexibility	Low	Medium	Very High
Impact Resistance	Low	Medium/High	High
Impact Absorption	Low	High	Very High
Heat Deflection Resistance	Low	High	Medium
Fatigue resistance under bending	Low	High	Very High
Glass transition temperature (T_g)	~333,1K	~378,15 K	~303,15 - 333,15 K

Table 1: Comparison of mechanical properties of PLA, ABS, and TPU polymers.

The final layout of the new drone is shown below.




Figure 1: IMAV 2025 Drone.

2.3. Simulation

To validate the structural design of the MAV, several simulations were carried out under different hovering and accelerated flight conditions using Fusion 360 software. However, due to space limitations, we will document only one simulation we consider the most relevant.

This static simulation was performed to analyze the aerodynamic forces and loads the drone is subjected to in flight. The material properties were obtained from the Composite Materials database as well as manufacturer specifications. Additionally, the simulation applied four loads of 30 N each at the motor positions, replicating the maximum thrust of the propellers, gravitational acceleration ($g = 9.807 \text{ m/s}^2$), and a linear acceleration of 2 m/s^2 , to ensure the integrity of the plate supporting the components. The results of the simulation are shown below.




Figure 2: Safety Factor




Figure 3: Displacement

The maximum Von Mises stress was 43.869 MPa, the 1st Principal stress was 45.859 MPa, and the 3rd

Principal stress was 17.845 MPa. Normal stresses were practically negligible, with the exception of Normal YY stress, which reached 45.381 MPa. In addition, the shear stresses XY, YZ, and XZ were nearly identical, at approximately 5.5 MPa. The table below shows the maximum displacement that the structure may undergo under these simulation conditions.

Direction	Displacement (m)
X	0.00094
Y	0.004764
Z	0.00094
Total	0.00507

Table 2: Displacement

Reaction force is the force exerted by a constraint (such as a fixed face or a bolt) to counteract applied loads and keep the model from moving. Based on the maximum reaction forces obtained (Table 3), the drone's fasteners are well-dimensioned for the expected stresses in each mission test, thus ensuring fixed parts remain secure and structural integrity is maintained.

Direction	Maximum reaction forces
X	103.996 N
Y	42.065 N
Z	97.302 N
Total	45.79 N

Table 3: Maximum reaction forces

Contact force is the actual load applied between components during a simulation, while contact pressure is the force per unit area acting along the contact surface. From the results in Table 4, it is observed that most stresses are transmitted to the landing gear. Therefore, contact forces and pressures will not cause physical damage to the drone's structure, even under extreme conditions.

Direction	Contact Force	Contact Pressure
X maximum	127.701 N	13.347 MPa
Y maximum	213.11 N	10.107 MPa

Z maximum	143.348 N	12.622 MPa
Total maximum	246.421 N	20.992 MPa

Table 4: Contact forces and pressures




Figure 4: Von Mises

In summary, the drone is well-dimensioned for extreme flight conditions and will meet the design requirements of the planned tasks. With a safety factor of 11.092, it is possible to reduce the diameter of the carbon fiber tubes; however, we opted to maintain them as they are to provide additional protection in case of crashes. Furthermore, stress distribution is well-balanced throughout the drone, ensuring the structural integrity of the frame.

2.4 Mechanical Solutions for the Tasks

The drone developed for IMAV 2025 has approximate dimensions of 0.48 m between motors diagonally and about 0.53 m considering the total structure, complying with the maximum limit of 0.7 m established by the indoor competition regulations. This configuration ensures adequate conditions for completing the 1×1 m gate crossing mission, providing a safety margin during flight. The combination of low-mass materials and high structural stiffness helps reduce oscillations in confined environments, preserving the stability required for autonomous navigation.

For the package delivery stage, the use of a mechanism inspired by the operating principle of the Petersen dredge is envisioned. The system operates through levers that move articulated paddles responsible for retaining and releasing the payload. Since the final shape of the package has not yet been established, the paddle design will be adjusted later to ensure compatibility with the dimensions of the transported object. The regulations specify the delivery of med-kits measuring 0.20 m × 0.10 m × 0.05 m and approximately 0.03 kg, which will guide the final sizing of the mechanism. This solution provides versatility and precision, allowing the package to remain stable during flight and ensuring controlled release at the designated point.

Additionally, to reduce vibration transmission to the optical system, the camera will be mounted on ABS supports fixed to the frame using screws incorporating elastomeric inserts. These elements, positioned between the support and the structure, act as decoupling points, absorbing part of the vibrations generated by the motors and propellers. This measure helps improve the quality of the captured images and increases the reliability of autonomous navigation in indoor environments.




Figure 5: Elastomeric inserts used between the camera mount

3 ELECTRICAL AND ELECTRONIC SYSTEM

In this section, the main elements that constitute the drone's electronic architecture are presented, emphasizing the balance between computational power, energy efficiency, and robustness, factors that are essential to ensure autonomy and reliability during the IMAV 2025 missions. Furthermore, the rationale behind the selection of each component is discussed, highlighting their alignment with the mission plan requirements.

3.1 NVIDIA Jetson Orin Nano Super

The choice of the onboard computer to be used for carrying out the tasks was based especially on its computational power and parallel processing capability for using artificial intelligence (AI) models. Considering these points, the chosen board was the NVIDIA Jetson Orin Nano Super Developer Kit, shown in figure 6, equipped with 8 GB of LPDDR5 Random Access Memory (RAM) [9-10].



Figure 6: NVIDIA Jetson Orin Nano Super Developer Kit

Considered as a compact edge AI computer, it uses a 6-core Arm Cortex-A78AE CPU (1.7 GHz) and an Ampere-architecture GPU with 1024 CUDA cores (32 Tensor cores), delivering up to 67 INT8 TOPS of AI performance. The board includes 8 GB of 128-bit LPDDR5 memory (102 GB/s) and supports expandable storage via microSD or NVMe SSD [9-10].

Its carrier board exposes 4× USB3.2 Gen2 Type-A ports, Gigabit Ethernet, DisplayPort, and dual M.2 slots. This computer runs on a wide voltage input range, from 9V to 19V DC, and it typically consumes around 7 to 15 W of power, but when fully loaded, it can consume up to 25 W. Considering this, it is powered via a battery-backed Ultimate Battery Elimination Circuit (UBEC) to ensure a stable input voltage, as detailed in figure 8.

3.2 Intel RealSense D435i

The Intel RealSense D435i, shown in figure 7, is a stereo depth camera with global-shutter sensors and an integrated 6-Degrees of Freedom (6-DoF) Inertial Measurement Unit (IMU)[11]. It produces depth maps at up to 1280×720 resolution at 30 fps, and an RGB video stream at 1920×1080@30 fps[12]. Its stereo baseline is approximately 50 mm, yielding a depth field of view of roughly 87°×58° (H×V). The useful range is about 0.3–3 m, with a minimum depth of 0.28 m at full resolution [13].



Figure 7: Intel RealSense D435i

The camera connects via USB-C 3.1 Gen 1 (5 Gbps), drawing its 5 V power directly from the host computer, simplifying its power requirements[13]. Therefore, the connection diagram for the use of the NVIDIA Jetson Orin Nano Super onboard computer and the Intel RealSense D435i camera is shown in figure 8.




Figure 8: Connection diagram between the NVIDIA Jetson Orin Nano and the Intel RealSense D435i camera
Source: Authorial

3.3 PIXHAWK 6C MINI

The Pixhawk 6C Mini[14], shown in Figure 9, was selected as the flight controller due to its outstanding performance in small- and medium-sized drones operating in constrained environments.

Previously, the team employed the Pixhawk 2.4.8, an older and outdated version of the controller, which lacked the advanced processing capabilities, sensor redundancy, and integration support for specific components.

The Pixhawk 6C Mini retains the core features of the standard Pixhawk 6C but stands out for its compact design, cost-effectiveness, and reduced weight, thereby enhancing its integration with components and peripherals essential for autonomous flight.



Figure 9: Pixhawk 6C Mini.

The Pixhawk 6C Mini is equipped with two processors: an FMU (STM32H743), a high-performance unit featuring an Arm® Cortex®-M7 core running at 480 MHz with 1 MB of SRAM, and an IO processor (STM32F103), based on a Cortex®-M3 core operating at 72 MHz with 64 KB of SRAM. In addition, it integrates onboard sensors that ensure precise navigation and control, including two accelerometer/gyroscope units (ICM-42688-P and BMI055), a magnetometer (IST8310), and a barometer (MS5611).

3.4 CAMERA IMX219

The IMX219 camera module, shown in Figure 10, features a 1/4" 8MP MIPI CSI-02 sensor, previously employed in the Raspberry Pi Camera Module V2 released in 2016. This camera provides native integration with the NVIDIA Jetson Orin Nano. The specifications of the IMX219 are listed in Table 5.

This module is sufficient to accomplish the mission requirements, both due to its ability to deliver high-resolution imaging at a considerable frame rate and its seamless integration with the onboard computer and other drone components.



Figure 10: Camera module IMX219.

Shutter Type	Rolling Shutter
Active Pixels	3280 (H) × 2464 (V)
Resolution	8MP
Image Sensor Format	1/4" type
Pixel Size	1.12μm×1.12μm
CSI-2 Data Output	2-lane Mode
Data Format	Raw Bayer 10bit

Table 5: IMX219 Camera Module Specifications

3.5 FRSKY Taranis X9 lite and FRSKY X8R receiver

The FrSky Taranis X9 Lite transmitter operates in the 2.4 GHz frequency band using the ACCESS protocol, also compatible with ACCST, ensuring reliable communication and strong resistance to interference. It supports up to 24 channels, is compatible with OpenTX/ErskyTX, and provides bidirectional telemetry, audible alarms, and haptic alerts. Powered by two 18650 batteries, it offers extended autonomy and low latency, ensuring fast command response. Its ergonomic design minimizes pilot fatigue during long flight sessions, enhancing operational comfort.

Thus, the integration of the Taranis X9 Lite results in a robust, ergonomic, and efficient radio system, suitable for both manual flights and autonomous missions. Figure 11 presents the FrSky Taranis X9 Lite transmitter with its main controls and interfaces.



Figure 11: FRSKY Taranis x9-Lite

3.6 Benewake serial TF - Luna

The Benewake TF-Luna[16] is a short-range LiDAR sensor based on Time of Flight (ToF) technology, designed to measure the distance to objects through the emission and reception of light pulses. It features a range of 0.2 to 8 m, a resolution of 1 cm, an accuracy of ±6 cm up to 3 m, and an adjustable update rate from 1 to 250 Hz,

ensuring fast and reliable responses. The sensor is compact ($35 \times 21.25 \times 12.5$ mm), lightweight (<5 g), and has low power consumption (<0.35 W), making it ideal for MAV integration.

In this project, it is mounted on the underside of the aircraft and used as the primary sensor for Z-axis positioning (altitude relative to the ground). Its selection is based on its high reliability, consistent readings, and robustness under various operating conditions, which are crucial for autonomous landings and low-altitude flights. Figure 12 shows the TF-Luna module, highlighting its compact form factor and communication interfaces.



Figure 12: Benewake TF-Luna LiDAR Sensor

3.8 Li-Po Battery

The Li-Po (Lithium Polymer) battery serves as the primary energy source for the system, supplying power to both the propulsion motors and electronic subsystems. Li-Po batteries are characterized by their high energy density and lightweight design, making them especially suitable for drones and other aerial robotics applications. The battery adopted in this project is a 4S (14.8V nominal) pack with a capacity of 5200 mAh. Its high discharge capability (C-rate) provides the necessary current to sustain stable flight operations and power-demanding maneuvers.



Figure 13: Lithium Polymer (Li-Po) battery pack (4S, 5200 mAh)

3.8 Motors and ESCs

To enable the MAV to operate within its six degrees of freedom and perform autonomous flights, a set of SunnySky X2212 motors was selected, combined with 9450 propellers. Performance simulations were carried out using the E-calc software to ensure adequate autonomy and lift capacity. The results indicated that the MAV is capable of flying for up to 14 minutes, generating

a total thrust of 3.6 kgf, with a maximum current consumption of 16.33 A per motor.

For motor speed control, four individual 45 A ESCs (LANRC model) were chosen, dimensioned based on the maximum current per motor specified in the SunnySky X2212 980 Kv datasheet, which is 16.33 A. By applying a safety factor of 2, it was determined that each ESC should support at least 32.66 A. Since the available commercial options were 35 A and 45 A, and the price difference between them was negligible, the version with the higher continuous operating current was selected.

4 SOFTWARE SYSTEM

The software system developed for IMAV 2025 was structured to meet the requirements of autonomy and robustness imposed by the indoor missions of the competition. Considering the absence of GPS signals, the presence of obstacles, and the variability of environmental conditions, it became necessary to combine different perception and control strategies into an integrated, high-performance architecture.

Therefore, the software was divided into two main axes.

The first axis focuses on navigation and positioning, addressing the problem of accurate and consistent vehicle localization in indoor environments through advanced GPU-accelerated Visual SLAM (vSLAM) techniques. This component was responsible for evolving limited visual-inertial odometry solutions, employed in previous years, into a complete mapping and loop closure pipeline capable of ensuring stability, resilience, and the generation of persistent environmental maps.

The second axis encompasses the detection modules, dedicated to visual recognition missions and interaction with the environment, such as navigating through aligned tunnel gates and landing on platforms under adverse conditions. For this purpose, vision models based on deep learning (YOLOv11) were employed, ensuring robust, reliable, and real-time detection even under variable lighting.

5.1 Navigation and Positioning

Autonomous UAV navigation in indoor environments presents specific challenges due to the absence of GPS signals and the variability of environmental conditions, such as illumination, texture, and geometric features. The team's development has gone through several iterations to address these challenges, culminating in a robust, GPU-accelerated vSLAM. This section details the technical progression from our previous stack to the new solution, focusing on hardware selection, software architecture, and performance considerations.

5.1.1 Visual-Inertial Odometry with Intel RealSense T265

The team's initial successful deployment for the IMAV 2023 competition utilized the Intel RealSense T265[17], a dedicated tracking camera that provides a 6-Degrees-of-Freedom (6-DoF) pose estimation. The T265 is an integrated, black-box visual-inertial odometry (VIO) system, not a full SLAM solution. Its core components include:

- **Dual Fisheye Cameras:** These provide a wide-angle stereo image pair.
- **Onboard Inertial Measurement Unit (IMU):** A low-power IMU provides high-frequency angular velocity and linear acceleration data.
- **Movidius Myriad 2 VPU:** This dedicated Vision Processing Unit performs all VIO computations on-device.

The T265's integrated VPU was a significant advantage, offloading computationally intensive tasks from the main flight computer (Raspberry Pi 4)[18]. The fusion of visual features and IMU data allowed for accurate and low-latency odometry, which was a key factor in our team's third-place finish.

The software stack was built on ROS 2 Humble [19], with the librealsense v2.50.0 [20] and realsense-ros 4.51.1 [21] packages providing the interface to the T265. The VIO data was published as an Odometry topic, which was then consumed by a custom vision_to_mavros [22] node. This node translated the T265's output into the standard MAVROS [23] Vision_Position_Estimate message, feeding the pose directly to the PixHawk flight controller. Despite its initial success, this setup presented several challenges:

1. **Vibration Sensitivity:** The T265's internal IMU was highly sensitive to the vibrations from the UAV's motors, leading to noise in the odometry output. While we implemented a 3D-printed case with vibration-dampening foam, it was a persistent issue that required careful mechanical tuning.
2. **Black-Box Limitation:** The T265's proprietary, closed-source pipeline prevented low-level parameter tuning or debugging. This was a critical limitation during the IMAV 2024 mission, where an unfortunate accident damaged the camera, and subsequent issues with stabilization were difficult to diagnose without insight into the sensor's internal state or recalibration solutions.
3. **No Global Consistency:** As a VIO system, the T265 estimates poses incrementally and does not perform loop closure. Over long trajectories, this led to cumulative drift, which was a concern for the larger competition arena.

5.1.2 Transition to a Full V-SLAM Solution with Isaac ROS

Based on our experiences, the team's objective for the 2025 competition is to develop a more robust, flexible, and feature-rich vSLAM solution. This required a complete overhaul of both our hardware and software stacks.

5.1.2.1 Hardware Transition and System Architecture

Feature	Previous Stack (IMAV 2023-2024)	New Stack (IMAV 2025)
Sensor	Intel RealSense T265	Intel RealSense D435i
Compute Platform	Raspberry Pi 4	NVIDIA Jetson Orin Nano Super (CUDA)
Localization Method	Visual-Inertial Odometry	Full Visual SLAM (vSLAM)
Core Software	Proprietary T265 firmware	NVIDIA Isaac ROS Visual SLAM
Primary Output	Odometry Pose	Global Pose + Occupancy Map
Key Limitation	Cumulative drift, black-box system	Computationally intensive, requires powerful hardware
Primary Advantage	Low power, minimal latency	Global pose consistency, map generation, flexibility

Table 6: Contact forces and pressures

The D435i is an active stereo camera with an IMU. Unlike the T265, it does not have an onboard VPU and instead provides raw sensor data streams (stereo IR images, RGB, depth, and IMU data). The new Jetson Orin Nano Super, is now responsible for processing this raw data, enabling us to run highly optimized, GPU-accelerated vSLAM algorithms. Beyond SLAM, the computational capacity of the Jetson also opens the door to deploying more demanding perception modules, such as deep-learning-based object detection and semantic segmentation, further extending the UAV's autonomous capabilities.

5.1.2.2 Software Pipeline: Isaac ROS Visual SLAM

The new software pipeline leverages NVIDIA's Isaac ROS framework, a collection of hardware-accelerated ROS 2 packages designed for robotics perception. Isaac ROS [24] is optimized to use

the Jetson's GPU, which is critical for real-time performance.

The new vSLAM stack with Isaac ROS follows a more sophisticated pipeline:

1. **Sensor Data Ingestion:** The `realsense-ros` package ingests synchronized stereo image pairs and IMU data from the D435i, publishing them as standard ROS 2 topics.
2. **Visual-Inertial Fusion:** The core Isaac ROS vSLAM node processes the raw sensor data. This stage combines visual features from the stereo cameras with the IMU measurements to produce a highly accurate VIO estimate.
3. **Map Creation and Pose Graph Optimization:** The vSLAM system builds a map of the environment as it explores. This map is a collection of keyframes and landmarks. As new keyframes are added, they are linked to previous ones in a pose graph. This graph is continuously optimized to minimize reprojection errors, which corrects for short-term drift.
4. **Loop Closure:** A key feature of vSLAM that VIO lacks. When the UAV revisits a previously mapped area, the system detects a "loop closure." This triggers a global optimization of the entire pose graph, effectively correcting all accumulated drift and creating a globally consistent map and trajectory. This functionality is crucial for achieving long-term localization accuracy in a large indoor arena.
5. **Occupancy Mapping:** An optional but critical component for the mission is the use of Isaac ROS `nvblox`. By subscribing to the vSLAM pose and D435i depth data, `nvblox` generates a dense 3D occupancy map of the environment. This map can be used for obstacle avoidance, path planning, and autonomous navigation in cluttered spaces.

5.1.3 Summary of Transition

The progression from a limited-capability VIO system to a full GPU-accelerated vSLAM pipeline represents a significant technological leap for our team. By transitioning from the Intel RealSense T265 to the D435i, and adopting the NVIDIA Jetson Orin Nano with Isaac ROS, we have developed a system that is not only more robust and accurate but also more flexible and capable. This new stack provides global pose consistency through loop closure, the ability to generate a persistent map of the environment, and a path for integrating future advanced perception algorithms. This enhanced system is the cornerstone of our autonomous navigation solution for the IMAV 2025 competition.

5.2 Detection

Computer vision has long been a key component of our UAV platform, as the primary means for object detection. Over the years, visual detection modules have been integrated to handle environment interaction and mission-specific challenges, evolving from simple handcrafted methods to modern deep learning-based approaches. In the IMAV 2025 context, these capabilities are especially relevant for tasks such as navigating through a tunnel of aligned gates and landing on moving platforms.

Onboard computer vision solutions for UAVs have historically employed approaches based on shape and color detection [25]. Classical methods use Hough transforms to identify lines and quadrilaterals, or threshold-based segmentation in color spaces such as HSV and YCbCr. These algorithms have the advantage of low computational cost, allowing real-time execution on simple embedded platforms. In micro aerial vehicle competitions, such as previous editions of IMAV, color filters were frequently applied to highlight gates painted in strong tones and contour detection was used to locate landing platforms [25].

Despite their simplicity, these shape- and color-based techniques suffer from critical limitations in adverse scenarios. Variations in natural or artificial lighting directly affect color segmentation, leading to false detections or complete loss of the target [26]. Environmental conditions such as smoke, dust, or shadows reduce edge contrast, compromising gradient- and contour-based detectors [27]. Furthermore, more complex structures, such as the hollow square gates aligned in a tunnel, generate multiple visual instances that confuse traditional geometric algorithms, making consistent identification of the correct target difficult. In the case of the landing platform, lateral movements and the presence of smoke increase visual variability to a point where classical filtering methods do not provide reliability [27].

These limitations justify the adoption of deep learning models for onboard perception. Convolutional detection networks, such as the YOLO family, demonstrate the ability to learn representations robust to lighting variations, partial occlusions, and perspective changes [28]. Combined with data augmentation techniques that simulate adverse conditions, these architectures become suitable for competitive and mission-critical environments [29]. The strategy proposed in this section relies on these capabilities: the use of YOLOv11 in a dual configuration for gate navigation, and YOLOv11 for landing on a platform, aims to overcome the inherent problems of approaches based solely on shapes and colors [25].

The proposed system will be structured into two main modules, each dedicated to one of the tasks defined in the IMAV 2025 competition: navigation through tunnels formed by square gates and landing on a solid platform. Both modules share the philosophy of using YOLO-family deep learning models, trained on dedicated datasets that simulate visual conditions similar to those encountered in the competition. Data augmentation techniques will be applied to include variations in lighting, noise, blur, and perspective changes, ensuring model robustness. Training will be conducted with parameter tuning of anchors and input dimensions to optimize real-time performance on embedded hardware.

For gate traversal, a look-ahead strategy composed of two specialized YOLOv11 models will be employed. The first will be trained exclusively to identify the frontal gate, while the second will be trained to detect the farthest gate in the sequence. During flight execution, the detections provided by the two models will be used by a PID controller configured to maintain the alignment of the distant gate relative to the frontal gate in the image. This logic induces the vehicle to remain centered along the tunnel axis, reducing abrupt corrections even in scenarios of structural misalignment. The Non-Maximum Suppression (NMS) stage will be fine-tuned to consolidate detection hypotheses and prioritize the most relevant opening in each frame.

In the landing task, a YOLOv11 model dedicated to detecting the solid square platform will be employed. The model is trained to handle variations in lighting, perspective, and platform motion, providing reliable real-time detections for precise positioning. Detection outputs will directly guide the descent controller, allowing the UAV to adjust its trajectory and align with the platform accurately. This approach relies solely on deep learning-based perception, ensuring safe and robust landings while meeting the real-time performance.

5 CONCLUSION

For IMAV 2025, the Black Bee Drones project integrated mechanical, electronic, and software solutions into a single platform. The new structural design, combining carbon fiber tubes, ABS-printed mounts, and TPU landing supports, resulted in a lightweight, rigid, and reliable frame, validated in simulation environments.

In terms of hardware, the adoption of the NVIDIA Jetson Orin Nano Super together with the Intel RealSense D435i provided the computational power required to execute GPU-accelerated Visual SLAM. This new pipeline, based on Isaac ROS, ensures global localization consistency through loop closure, a critical feature for

long-duration missions in GPS-denied environments, such as indoor scenarios.

Furthermore, the team's software approach incorporated advanced computer vision with YOLOv11, enabling robust real-time object detection. This intelligent perception system allows the drone to perform tasks such as navigating through tunnel-like aligned gates and autonomously landing on moving platforms, surpassing traditional methods based solely on color or shape.

Therefore, the Black Bee Drones team presents a prototype that combines reliability, performance, and technologies, capable of overcoming challenges faced in previous projects while driving continuous progress in research and development within the field of unmanned aerial vehicles.

REFERENCES

- [1] KUMAR, N.; MISHRA, S. *Carbon Fiber in the Aerospace Industry*. ResearchGate, 2022. Disponível em: https://www.researchgate.net/publication/394082670_Carbon_Fiber_in_the_Aerospace_Industry. Acesso em: 27 ago. 2025.
- [2] RAHMAN, M. *Thermal Analysis of ABS*. COMSOL Conference, 2019. Disponível em: https://www.comsol.jp/paper/download/361301/rahman_paper.pdf. Acesso em: 27 ago. 2025.
- [3] WEVOLVER. *Understanding ABS Heat Resistance: A Comprehensive Guide*. Wevolver, 2024. Disponível em: <https://www.wevolver.com/article/understanding-abs-heat-resistance-a-comprehensive-guide>. Acesso em: 27 ago. 2025.
- [4] 3DTRCEK. *ABS Filament Temperature Resistance: A Comprehensive Guide*. 3DTRCEK, 2023. Disponível em: <https://3dtrcek.com/en/blog/post/abs-filament-temperature-resistance-a-comprehensive-guide>. Acesso em: 27 ago. 2025.
- [5] SUNLU. *The Properties of TPU Material and Its Application in 3D Printing*. Sunlu, 2024. Disponível em: <https://www.sunlu.com/blogs/products-knowledge/the-properties-of-tpu-material-and-its-application-in-3d-printing>. Acesso em: 27 ago. 2025.
- [6] XOMETRY. *3D Printing Flexible Materials*. Xometry Pro, 2024. Disponível em: <https://xometry.pro/en/articles/3d-printing-flexible-materials/>. Acesso em: 27 ago. 2025.
- [7] ZHANG, Y. et al. Polyurethane Elastomers for Biomedical Applications: Mechanical Properties and Fatigue Resistance. *Polymers*, v. 11, n. 4, p. 694, 2019. Disponível em: <https://pmc.ncbi.nlm.nih.gov/articles/PMC6432304/>. Acesso em: 27 ago. 2025.
- [8] AHMAD, R. et al. *Additive Manufacturing of Thermoplastic Polyurethane Elastomers: A Review*. ResearchGate, 2023. Disponível em:

https://www.researchgate.net/publication/384786975_Adaptive_Manufacturing_of_Thermoplastic_Polyurethane_Elastomers_A_Review. Acesso em: 27 ago. 2025.

[9] NVIDIA. (n.d.). Jetson Orin Nano Super Developer Kit. Retrieved September 4, 2025, from <<https://www.nvidia.com/en-us/autonomous-machines/embedded-systems/jetson-orin/nano-super-developer-kit/>>

[10] Sseed Studio. (n.d.). NVIDIA® Jetson Orin™ Nano Developer Kit. Retrieved September 4, 2025, from <<https://www.sseedstudio.com/NVIDIA-Jetson-Orin-Nano-Developer-Kit-p-5617.html>>

[11] FRAMOS. (n.d.). Depth Camera D435I (bulk). Retrieved September 4, 2025, from <<https://framos.com/products/3d/3d-cameras/depth-camera-d435i-bulk-22610/>>

[12] B&H Photo Video. (n.d.). Intel RealSense Depth Camera D435I. Retrieved September 4, 2025, from <https://www.bhphotovideo.com/c/product/1495418-REG/intel_82635d435idk5p_realsense_depth_camera_d435i.html>

[13] RealsenseAI Store. (n.d.). Buy Intel RealSense Depth Camera D435I. Retrieved September 4, 2025, from <<https://store.realsenseai.com/buy-intel-realsense-depth-camera-d435i.html>>

[14] PX4 Autopilot. *Pixhawk 6C Mini*. Available at: https://docs.px4.io/main/en/flight_controller/pixhawk6c_mini. Accessed on: August 29, 2025.

[15] Arducam. *IMX219 Native Camera Module for Nvidia Jetson*. Available at: <https://docs.arducam.com/Nvidia-Jetson-Camera/Native-Camera/imx219/#products>. Accessed on: August 29, 2025.

[16] Benewake (Beijing) Co., Ltd., *TF-Luna LiDAR Module (Short-range Distance Sensor) Datasheet*, 2024. Available:

<https://en.benewake.com/uploadfiles/2024/04/20240426135921367.pdf>. Accessed 29 August 2025.

[17] Intel Corporation, "Intel® RealSense™ Tracking Camera T265 Datasheet," 2019. [Online]. Available: <https://www.intel.com/content/dam/support/us/en/documents/emerging-technologies/intel-realsense-technology/IntelRealSenseTrackingT265Datasheet.pdf>. Accessed: Sep. 5, 2025.

[18] Raspberry Pi Ltd., "Raspberry Pi 4 Model B Datasheet," 2024. [Online]. Available: <https://datasheets.raspberrypi.com/rpi4/raspberry-pi-4-datasheet.pdf>. Accessed: 5 set. 2025.

[19] Open Robotics, "ROS 2 Humble Hawksbill," 2025. [Online]. Available: <https://docs.ros.org/en/humble/Releases.html>. Accessed: 5 set. 2025.

[20] Intel RealSense, "Intel RealSense SDK 2.50.0," 2021. [Online]. Available: <https://github.com/IntelRealSense/librealsense/releases/tag/v2.50.0>. Accessed: 5 set. 2025.

[21] Intel RealSense, "realsense-ros 4.51.1," 2022. [Online]. Available: <https://github.com/IntelRealSense/realsense-ros/releases/tag/4.51.1>. Accessed: 5 set. 2025.

[22] Black-Bee-Drones, "vision_to_mavros," GitHub repository, 2025. [Online]. Available: https://github.com/Black-Bee-Drones/vision_to_mavros. Accessed: 5 set. 2025.

[23] MAVLink, "MAVROS 2.10.1," 2024. [Online]. Available: <https://github.com/mavlink/mavros/releases/tag/2.10.1>. Accessed: 5 set. 2025.

[24] NVIDIA, "Isaac ROS," 2025. [Online]. Available: <https://nvidia-isaac-ros.github.io/releases/index.html>. Accessed: 5 set. 2025.

[25] Van Garvanov, Georgi Tsonkov, and Magdalena Garvanova. Objects detection in drone swarm surveillance by color features. In 2025 19th Conference on Electrical Machines, Drives and Power Systems (ELMA), pages 1–5, 2025.

[26] Hiraj Gandhi, Lerrel Pinto, and Abhinav Gupta. Learn-ing to fly by crashing. In 2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), pages 3948–3955, 2017.

[27] Mario Bijelic, Tobias Gruber, Fahim Mannan, Florian Kraus, Werner Ritter, Klaus Dietmayer, and Felix Heide. Seeing through fog without seeing fog: Deep multimodal sensor fusion in unseen adverse weather. In 2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), pages 11679–11689, 2020.

[28] Haoran Chen, Jingyu Wang, Jingwei Li, Yu Qiu, and Dengyin Zhang. Small object detection for drone image based on advanced yolov7. In 2023 42nd Chinese Control Conference (CCC), pages 7453–7458, 2023.

[29] Bawei Hou, Bo Li, Yang Li, and Jiao Ma. Drone aerial object detection based on improved yolov8. In 2025 7th International Conference on Natural Language Processing (ICNLP), pages 448–452, 2025.