Feasibility of Onboard Vision-Based Detect-and-Avoid for Small UAS Under Size, Weight, and Power Constraints

Haoran Xu[1] **and Liying Chen**[2]

[1] ***Nanjing Institute of Mechanical Engineering, Department of Mechanical Engineering, 112 Longpan Road, Xuanwu District, Nanjing 210042, China***
[2] ***Chengdu University of Technology and Design, Department of Mechanical Engineering, 56 Tianfu Avenue, Wuhou District, Chengdu 610041, China***

Abstract

Small unmanned aircraft systems increasingly share low-altitude airspace with crewed aircraft, rotorcraft, and other unmanned platforms whose trajectories are only partially observable and weakly cooperative. In this environment, onboard detect-and-avoid capabilities are necessary to support beyond visual line of sight operations when reliance on ground-based surveillance or cooperative transponders is not assured. Vision-based sensing is attractive for small platforms because of its geometric richness and the availability of compact low-cost cameras, but its feasibility is constrained by severe limits on mass, power, computation, and thermal dissipation. This paper examines the feasibility of onboard vision-based detect-and-avoid for small unmanned aircraft systems subject to realistic size, weight, and power constraints by integrating sensor modeling, algorithmic complexity analysis, and closed-loop encounter-level performance. The discussion focuses on monocular and stereo visible-band configurations mounted on multirotor and fixed-wing platforms with maximum take-off mass below 25 kg and continuous electrical power budgets below 80 W. A dynamical engagement framework is used qualitatively to relate detection range distributions, track continuity, and decision latency to miss-distance statistics without presupposing a particular regulatory standard of performance. The analysis highlights the coupled role of optics, pixel-level signal-to-noise, embedded inference latency, and maneuver authority in shaping achievable detect-and-avoid envelopes, while acknowledging that environmental variability and non-cooperative traffic behavior introduce uncertainties that limit deterministic guarantees. The results collectively indicate conditions under which onboard vision sensing constitutes a technically viable component of layered separation assurance architectures for small unmanned aircraft.

1 Introduction

Modern low-altitude airspace increasingly accommodates diverse aerial actors that range from lightweight multirotor platforms conducting infrastructure inspection to fixed-wing small unmanned aircraft performing logistics, environmental monitoring, and public safety missions [1]. These operations occur within a limited vertical envelope, typically below 400 feet above ground level, where the interaction between crewed general aviation traffic, obstacles, and unmanned systems becomes dense and complex. The multiplicity of flight purposes and operational patterns introduces heterogeneity in flight trajectories, velocity distributions, and maneuvering behaviors. This complexity challenges existing separation assurance concepts developed for traditional air traffic management, which presume cooperative surveillance, standardized equipment, and centralized control. Small unmanned aircraft, by contrast, often operate with decentralized command architectures and without transponders, depending instead on local sensing and autonomy to ensure safety.

A central difficulty lies in the limited feasibility of universal cooperative surveillance. Many small platforms cannot carry certified transponders or automatic dependent surveillance broadcast units because of mass, power, and cost limitations. Similarly, ground-based radar or networked sensor coverage may not extend into the low-altitude corridors where unmanned aircraft fly, particularly in rural and semi-urban settings [2]. Even where network links exist, bandwidth and latency constraints restrict the utility of continuous uplink or downlink for real-time conflict detection. Therefore, it cannot be assumed that all relevant traffic participants are visible to any centralized system. The resulting partial observability of the airspace motivates research into onboard detect-and-avoid systems capable of sensing, interpreting, and responding to traffic autonomously.

For such onboard systems, perception of nearby aircraft or obstacles is not a luxury but a safety-critical requirement. Detect-and-avoid capability encompasses three linked functions: detection of intruders within a defined volume around the ownship, estimation of their trajectories relative to the ownship, and timely generation of maneuver commands to prevent loss of separation. Implementation of these functions on small unmanned aircraft must respect strict size, weight, and power constraints. Each gram of additional mass affects endurance, and each watt of power devoted to sensing or computation reduces energy available for propulsion [3]. Embedded processors must dissipate heat without active cooling, and the physical volume available for cameras or sensors is bounded by aerodynamic and structural limitations. The detect-and-avoid subsystem must therefore achieve sufficient sensing performance using minimal hardware, a demanding engineering problem given the range and reliability requirements associated with air collision avoidance.

Vision-based sensing emerges as a particularly attractive solution because cameras offer high spatial resolution per unit mass and can function passively without emitting energy. Passive vision aligns naturally with the power and weight budgets of small unmanned aircraft. The same small optical systems used for navigation or mapping can be leveraged for detect-and-avoid, thereby promoting hardware reuse and integration. However, vision-based systems introduce dependencies on environmental factors that are difficult to control. Illumination levels, weather conditions, and background textures exert substantial influence on detection performance [4]. Bright sun angles can cause glare, while overcast skies may reduce contrast between an aircraft and its background. Fog, haze, or rain degrade visibility, while cluttered ground backgrounds or moving shadows increase false alarm likelihood. Thus, the reliability of a purely vision-based system fluctuates with external conditions.

Despite these limitations, the ability of modern computer vision algorithms to extract geometric and semantic information from images has advanced significantly. Detection of small, moving objects against dynamic backgrounds has become feasible even on embedded processors, though still computationally demanding. The typical approach involves detecting potential targets, tracking them over time, and analyzing motion patterns to estimate whether they represent collision threats. Unlike radar, which directly provides range and velocity, cameras yield only bearing and bearing rate information. Range estimation therefore relies on motion parallax, apparent size variation, or stereo disparity [5]. Each of these methods has sensitivity limits determined by optics and image quality. Small targets at large distances subtend only a few pixels, making reliable detection challenging unless noise levels are extremely low and exposure control is precise.

The absence of direct range information distinguishes vision-based detect-and-avoid from radar-based systems. Range must be inferred indirectly from sequences of images, often through filtering or probabilistic estimation. For example, if the bearing of a detected object remains approximately constant over successive frames while its apparent size increases, a potential collision trajectory is implied. This "constant-bearing, decreasing-range" heuristic has been used in both biological and robotic contexts as an indicator of impending collision. Implementing it robustly on a resource-limited embedded system requires accurate stabilization, timing, and image registration to separate genuine bearing constancy from apparent effects due to platform vibration or camera motion. [6]

Autonomy in detect-and-avoid demands that perception, estimation, and control functions interact seamlessly. Once an intruder is detected, the system must estimate closure time and feasible avoidance trajectories consistent with flight dynamics. For multirotor aircraft, rapid lateral or vertical maneuvers may be possible, whereas fixed-wing aircraft are constrained by stall limits and minimum turning radii. The decision logic must thus be parameterized by platform characteristics. It must also account for uncertainties in detection and estimation; premature or spurious avoidance actions can deplete energy reserves and disrupt missions, while delayed reactions risk collision. Balancing these opposing risks within the computational and sensing constraints of small unmanned aircraft represents a central feasibility challenge.

The detect-and-avoid problem can be considered at two timescales [7]. At the strategic level, the aircraft plans its trajectory to avoid entering conflict zones predicted over horizons of tens of seconds or minutes. At the tactical level, it responds to imminent threats within seconds. Small unmanned aircraft with limited sensing range and computational resources will necessarily rely more heavily on tactical avoidance, because their sensors may not detect distant intruders early enough for long-term trajectory planning. Vision-based systems are particularly suited to the tactical regime, where relative motion cues and short-range detections dominate. However, achieving the necessary update rates and processing speeds requires careful optimization of both hardware and algorithms. Delays of even a few hundred milliseconds can materially reduce the time available for avoidance, especially at closure speeds exceeding 30 m/s.

Another fundamental aspect concerns field of view. Cameras have finite angular coverage, and obstructions from airframe structures, rotors, or payloads can produce blind zones [8]. Multiple cameras may be required to achieve near-spherical coverage, but each additional unit adds mass and power consumption. Furthermore, wider fields of view obtained through fisheye lenses reduce angular resolution and hence detection range. Designers must therefore trade between coverage and sensitivity, considering the typical geometries of encounters in the intended operational domain. For example, head-on encounters may dominate in certain flight corridors, favoring forward-facing cameras with moderate fields of view and high resolution. In contrast, urban operations with potential intrusions from multiple directions might justify panoramic configurations despite their higher resource cost.

Thermal management is another constraint seldom highlighted in high-level detect-and-avoid discussions but critical in practice. Embedded processors executing intensive computer vision workloads generate heat that must be dissipated through convection and conduction [9]. In thin air at altitude or enclosed fuselage spaces with limited airflow, temperature rise can lead to throttling or failure. Passive heat sinks and careful placement relative to cooling airflow become essential design considerations. Yet these add to mass and volume, illustrating again how each physical parameter interacts with others in the size, weight, and power design space.

The interplay of these constraints defines an achievable region of system performance rather than a single operating point. Within this region, performance metrics such as detection probability versus range, false alarm rate, processing latency, and power draw are interdependent. Improving one metric often degrades another. Increasing frame rate enhances temporal resolution but raises processing load and power consumption [10]. Increasing exposure time improves sensitivity but introduces motion blur. Employing deep neural networks can improve detection robustness but may exceed the capabilities of available hardware. Consequently, feasibility assessments must evaluate system performance not only in isolation but also in the context of the complete aircraft energy and control budget.

From an operational perspective, the detect-and-avoid systems purpose is to support safe and predictable integration of unmanned aircraft into shared airspace. This goal imposes implicit reliability expectations. Even if regulations do not mandate quantitative probabilities of midair collision avoidance for small unmanned aircraft, public acceptance and insurance considerations will demand consistent performance. Vision-based systems, given their environmental dependencies, will likely need to demonstrate reliability through extensive statistical evidence collected over diverse conditions [11]. This requirement interacts with the hardware constraints because testing under variable illumination, clutter, and weather may expose performance degradations that can only be mitigated by hardware improvements, such as higher dynamic range sensors or additional cameras. Hence, the process of validating feasibility becomes iterative between algorithmic optimization and hardware adaptation.

Another aspect that affects practical feasibility is integration with other onboard functions. Cameras used for detect-and-avoid may also serve navigation, mapping, or inspection purposes, sharing bandwidth and computation with these functions. Time-multiplexing or prioritization schemes are required to ensure that safety-critical detect-and-avoid processing preempts less critical tasks during potential conflicts. This introduces complexity into software architecture and scheduling, as well as implications for safety assurance. The detect-and-avoid function must remain responsive even under heavy system load or degraded conditions. Techniques such as fixed-time scheduling, watchdog monitoring, and fault-tolerant processing can mitigate these risks but consume additional resources. [12]

The dynamics of small unmanned aircraft themselves contribute uncertainty. Lightweight multirotors are susceptible to wind gusts, which can cause abrupt attitude changes and vibrations. These motions induce image jitter that complicates small-target detection, especially at longer ranges where targets occupy few pixels. Image stabilization and inertial compensation can reduce these effects but require additional sensors and computational effort. Similarly, the vibratory environment can affect camera calibration and focus, leading to degraded image quality over time. Sustaining reliable detect-and-avoid performance under such conditions requires robust mechanical design as well as algorithmic compensation.

Future developments in imaging technology may alleviate some of these constraints [13]. High-sensitivity, low-noise image sensors capable of operating across wider dynamic ranges could extend reliable detection to lower illumination levels without increasing power consumption. Advances in neuromorphic or event-based vision sensors, which report only pixel changes rather than full frames, offer potential for significant power and bandwidth savings. Such sensors align naturally with the motion-centric nature of detect-and-avoid tasks, although their maturity for safety-critical applications remains limited. Integrating these technologies into practical systems will still require adherence to the overall size, weight, and power envelope dictated by small unmanned aircraft design.

In parallel, progress in low-power computing platforms continues to expand feasible algorithmic options. Specialized accelerators for convolutional neural networks and dedicated vision processors can deliver high performance per watt, narrowing the gap between research-grade algorithms and embedded implementation. Nevertheless, total available power on small platforms remains finite, and even efficient processors add mass and complexity [14]. Consequently, optimization must occur at the system level, balancing computation, sensing, and control rather than optimizing any single component in isolation.

Modern low-altitude airspace presents a challenging environment in which small unmanned aircraft must operate safely among a variety of cooperative and non-cooperative participants. The detect-and-avoid capability is essential for achieving this safety when external infrastructure and universal equipage cannot be relied upon. Vision-based sensing offers an appealing route due to its passive nature and compatibility with small, power-limited platforms. However, the dependence on environmental conditions, lack of direct range measurement, and stringent hardware constraints combine to make implementation complex. Achieving reliable detect-and-avoid performance requires careful co-design of optical systems, algorithms, and platform integration, supported by rigorous testing across operational conditions. The feasibility of such systems rests not on any single technological breakthrough but on balanced engineering trade-offs that reconcile sensing and computation with the physical realities of small unmanned aircraft flight.

The feasibility question is therefore not only whether vision-based detect-and-avoid can be made to work at all, but whether it can be made to work consistently enough within realistic hardware budgets and operating scenarios to warrant adoption as a substantive contributor to risk reduction [15]. This question is complicated by the absence of universally accepted quantitative performance metrics for small unmanned aircraft detect-and-avoid and by the diversity of potential operations, from structured corridors with constrained traffic patterns to unstructured environments with heterogeneous vehicles and pilot behaviors. A neutral assessment must focus on the internal coherence of sensing and decision processes, the explicit accounting of size, weight, and power limitations, and the sensitivity of outcomes to environmental and behavioral uncertainties, rather than on optimistic extrapolation from isolated demonstrations.

This paper develops such an assessment by combining physical models of camera sensing, embedded computation, and aircraft dynamics with abstracted encounter models to relate hardware-level and algorithm-level design choices to probabilities of timely conflict detection and avoidance. The intention is not to prescribe detailed certification criteria, but to examine the structural relationships that determine when onboard vision-based detect-and-avoid can function as one element of a layered safety architecture for small unmanned aircraft, and when its limitations are likely to dominate.

Beyond technical challenges, the detect-and-avoid function interacts with communication, navigation, and surveillance architectures that may include cooperative transponders, cellular or satellite links, and ground-based sensing. For small unmanned aircraft that cannot rely on redundant avionics or extensive certification pedigree, it is reasonable to expect that onboard vision-based detect-and-avoid will operate alongside, rather than in isolation from, these other sources of information where available. This motivates formulations in which vision contributes bearing-only or bearing-plus-classification information into a fusion framework that can also incorporate cooperative data when present [16]. Under size, weight, and power constraints, however, fusion must be implemented in modest hardware without assuming the availability of full-size airborne surveillance systems, implying that the marginal benefit of vision must be evaluated relative to its incremental resource cost.

The introduction of autonomy in detect-and-avoid decision-making also raises questions of predictability and explainability. For systems that may rely on learning-based components trained on large image corpora, systematic characterization of failure modes is challenging. From a feasibility standpoint, this difficulty does not by itself preclude deployment but signals the importance of binding algorithm performance to measurable quantities such as detection range statistics and false alarm rates over relevant data sets. The present analysis frames feasibility in such measurable terms, avoiding reliance on unquantified generalization claims and instead tracing the implications of observed or modeled performance through to encounter-level outcomes under clearly stated assumptions.

Table 1. Representative Size, Weight, and Power Budgets for Small UAS Platforms (Total width $\leq$ 12 cm)

Platform Type	Max Takeoff Mass (kg)	Payload Capacity (g)	Power Budget (W)	Detect-and-Avoid Allocation (W)
Micro Multirotor	1.5	200	80	8
Small Multirotor	5.0	800	200	12
Fixed-Wing (Short Range)	10.0	1500	300	15
Fixed-Wing (Long Endurance)	20.0	2500	400	20

Table 2. Camera Configurations and Optical Parameters for Vision-Based Detect-and-Avoid (Total width $\leq$ 12 cm)

Camera Type	Resolution (pixels)	Focal Length (mm)	Field of View (degrees)	Power Consumption (W)
Monocular CMOS	1920×1080	6.0	78	2.0
Stereo Pair CMOS	1280×720 (each)	4.0	90	4.5
Global-Shutter CMOS	2048×1536	8.0	70	3.2
Event-Based Sensor	640×480	3.0	100	1.2

2 System-Level Constraints and Operational Scenarios

Size, weight, and power constraints for small unmanned aircraft systems arise from aerodynamic, structural, and energy storage considerations that leave limited margin for mission-specific payloads. A typical electric multirotor operating with a maximum take-off mass of a few kilograms may allocate less than several hundred grams and a small fraction of its power budget to all auxiliary payloads, including sensing and computation [17]. Fixed-wing platforms optimized for endurance often exhibit higher payload fractions but must still reconcile mass addition with impacts on wing loading, stall speed, and climb per-

Table 3. Embedded Processing Platforms and Performance Characteristics (Total width ≤ 12 cm)

Processor Model	Peak Compute (GFLOPS)	Typical Power (W)	Thermal Limit (řC)	Memory Bandwidth (GB/s)
ARM Cortex-A72	25	5	80	6
NVIDIA Jetson Nano	500	10	85	25
Qualcomm QRB5165	1000	15	95	30
Custom FPGA Board	150	8	70	12

Table 4. Environmental Factors Affecting Detection Performance (Total width ≤ 12 cm)

Condition	Illumination (lux)	Visibility Range (km)	Relative Detection Range Reduction (%)	False Alarm Rate Increase (%)
Clear Daylight	10000	10	0	0
Overcast	3000	8	15	8
Hazy	1000	4	35	20
Foggy	300	1	60	40
Low Sun Angle	8000	6	25	15

Table 5. Detection Performance Metrics as a Function of Range (Total width ≤ 12 cm)

Range (m)	Detection Probability (%)	False Alarm Probability (%)	Processing Latency (ms)	Confidence Score (01)
100	98	1	50	0.95
200	85	2	65	0.88
300	70	3	80	0.76
400	45	5	100	0.58
500	25	6	120	0.43

formance. In both cases, integration of detect-and-avoid hardware must not materially degrade the stability, controllability, or endurance required for the intended mission. This implies that camera modules, lenses, enclosures, interconnects, and computing boards must be selected with attention to both their individual properties and their combined effects on mass distribution and aerodynamic drag.

Power constraints are closely tied to propulsion system design. Battery-powered multirotors may allocate on the order of 100 W to 400 W for propulsion during hover and climb, with limited additional headroom available. Embedded processors suitable for onboard vision, depending on architecture and operating point, may consume from a few watts to tens of watts. Even modest additional consumption can shorten mission duration or reduce thrust margin, and thermal dissipation for such components is non-trivial on compact airframes with limited airflow. In practice, design margins are allocated for worst-case conditions, including high ambient temperature and low air density, leading to conservative limits on continuous processing power [18]. These constraints influence

Table 6. Processing Latency Contributions by Pipeline Stage (Total width $\leq$ 12 cm)

Stage	Mean Latency (ms)	Power Share (%)	Memory Use (MB)	Comments
Image Acquisition	10	15	30	Exposure and transfer time
Preprocessing	20	20	60	Filtering, normalization
Object Detection	35	35	200	Neural inference or feature extraction
Tracking	15	20	80	Kalman or particle filtering
Decision Logic	10	10	20	Threat classification and command output

Table 7. Comparison of Vision-Based and Non-Vision Detect-and-Avoid Modalities (Total width $\leq$ 12 cm)

Modality	Range (m)	Power Use (W)	Mass (g)	Cooperative Requirement
Vision (Visible Band)	300600	8	120	None
Stereo Vision	200400	10	200	None
ADS-B Receiver	>2000	3	50	Yes
Small Radar	8001200	20	500	None
Acoustic Array	150300	5	300	None

Table 8. Representative Avoidance Performance as Function of Closure Rate (Total width $\leq$ 12 cm)

Closure Rate (m/s)	Required Warning Time (s)	Required Detection Range (m)	Probability of Avoidance (%)	Platform Maneuver Capability (m/s^2)
20	2.0	40	99	5.0
40	3.0	120	92	4.0
60	4.0	240	80	3.0
80	5.0	400	65	2.5
100	6.0	600	45	2.0

feasible frame rates, model sizes, and numbers of cameras.

Mechanical integration constraints further restrict the placement and orientation of vision sensors. Cameras require unobstructed fields of view free from propeller blades,

Table 9. Power Distribution and Allocation in Small UAS (Total width $\leq$ 12 cm)

Subsystem	Power Use (W)	Power Share (%)	Operating Voltage (V)	Typical Duty Cycle (%)
Propulsion	150	75	14.8	100
Avionics & Control	15	7	5.0	100
Navigation Sensors	8	4	5.0	100
Detect-and-Avoid System	10	5	5.0	90
Communications	12	6	5.0	60

Table 10. Synthetic Evaluation Summary for Vision-Based Detect-and-Avoid (Total width $\leq$ 12 cm)

Scenario Type	Mean Detection Range (m)	Latency (ms)	False Alarms / hr	Missed Detections (%)
Clear Sky Head-On	450	70	2.1	3.0
Crossing Path (Hazy)	250	85	4.0	8.5
Overcast Opposite Heading	300	90	3.2	5.0
Partial Occlusion Urban	200	120	5.5	11.0
High-Clutter Terrain	180	130	6.0	14.0

landing gear, and payload structures that could introduce partial occlusions or stray reflections. Mounting on booms or extended structures can reduce occlusion but introduces structural loads and vibrational modes that may blur images or complicate stabilization. Enclosures must protect against dust, moisture, and impact without adding excessive mass. Cabling and connectors must maintain signal integrity and power delivery while resisting fatigue. All of these aspects form part of the effective size and weight budget of the detect-and-avoid system and interact with power dissipation by altering local airflow. [19]

Operational scenarios determine the class of encounters that the detect-and-avoid system must address. For small unmanned aircraft operating at low altitudes in sparsely populated areas, most conflicts may involve other small unmanned platforms or occasional general aviation aircraft transiting at moderate speeds. In suburban or peri-urban environments, interactions with helicopters, light aircraft, and elevated structures become more prevalent. In dense urban settings, line-of-sight to potential intruders may be obstructed by buildings, and background clutter becomes highly structured. Each scenario is associated with characteristic closure rates, relative altitudes, angular distributions of potential intruders, and background radiance patterns. A system optimized for detecting small dark targets against bright sky may exhibit strong performance in rural head-on encounters but reduced performance when faced with complex urban textures and partial occlusion.

Weather and illumination conditions complicate this picture [20]. Low sun angles create glare and elongated shadows, overcast conditions reduce contrast, and haze attenuates small features at distance. Precipitation and fog can severely limit effective vision range, independent of hardware capability. Small unmanned aircraft are likely to be operated in a wide variety of such conditions, either by design or due to rapidly changing environments. Since detect-and-avoid systems cannot rely on controlled laboratory conditions, a feasibility assessment must consider not only best-case performance but also degradation patterns across realistic operating envelopes. From a pragmatic standpoint, operators may impose restrictions that avoid the most challenging conditions; however, reliance on such restrictions should be explicit rather than implicit in feasibility arguments.

In this context, detect-and-avoid feasibility is closely linked to achievable coverage and warning times. Sensor placement and number determine the angular region in which intruders may be detected without severe foreshortening or glare, while platform agility determines the amount of time required to maneuver out of potential conflicts once they are identified [21]. Small multirotor systems with high thrust-to-weight ratios can produce rapid lateral or vertical accelerations, allowing effective avoidance maneuvers even when detection occurs at modest ranges. Fixed-wing systems with higher forward speeds and limited normal-load capability require longer lead times. Any claim of feasibility for an onboard vision-based system must therefore be tied to the specific combination of platform performance, sensor configuration, and operational constraints in question.

Quantitatively, representative small unmanned aircraft may reserve on the order of 5% to 15% of their total electrical power budget for payloads beyond essential flight control, leaving perhaps 5 W to 20 W for sensing, onboard processing, and communications combined. Within this envelope, continuous operation of a vision-based detect-and-avoid pipeline must coexist with navigation computation and any mission-specific processing. Mass allocations may similarly be limited to 5% to 10% of total take-off mass, and a significant portion of that can already be consumed by primary mission payloads such as specialized cameras or delivery mechanisms. The detect-and-avoid subsystem therefore may be constrained to a few tens of grams for cameras and cables and a similar order for processing hardware. Thermal design must ensure that this hardware can reject heat to the environment without dedicated active cooling, which is rarely feasible on such platforms. [22]

Electromagnetic compatibility adds another subtle dimension. High-speed digital interfaces for cameras and processors can produce emissions that couple into communication systems or navigation sensors. Shielding and filtering measures to mitigate such coupling may add mass and volume, further tightening size and weight margins. The practical implication is that detect-and-avoid integration cannot be considered solely from the standpoint of individual component specifications; instead, system-level constraints encompassing power, mass, aerodynamics, thermal behavior, and electromagnetic compatibility jointly define the feasible region for onboard vision hardware.

3 Vision Sensing Architectures for Small UAS Detect-and-Avoid

Onboard vision-based detect-and-avoid architectures can be decomposed into sensor frontend, processing backend, and decision logic, each subject to size, weight, and power constraints. The sensor frontend encompasses camera modules, optics, and mechanical mounting. For small unmanned aircraft, cameras with small-format sensors and wide-angle lenses are frequently considered to maximize field of view per device [23]. However, wide-angle optics reduce angular resolution per pixel for distant targets, effectively lowering maximum detection range for a given target size and contrast. Narrower fields of view improve range but require multiple cameras or gimballed mounts to achieve coverage, increasing complexity and mass. Selection of focal length and sensor resolution therefore reflects a trade between instantaneous coverage and detection range that must be balanced in light of both encounter statistics and processing capacity.

Spectral sensitivity influences robustness. Visible-band sensors are efficient, low-cost, and compact, but their performance is coupled to ambient illumination and can be challenged by low sun or high dynamic range scenes. Near-infrared sensitivity can improve contrast for some materials and mitigate haze effects, yet may require specialized optics and introduce sensitivity to thermal emission. Long-wave infrared offers night capability and thermal contrast but is typically constrained by higher sensor cost, more demanding calibration, and significant power consumption [24]. For small unmanned aircraft with constrained power budgets, extensive cooling or large-format infrared arrays may be infeasible, suggesting that visible or near-infrared solutions with careful algorithm design are more compatible with size, weight, and power limitations.

Processing backend architectures must ingest image streams, perform detection and tracking, and feed outputs to guidance and control. Centralized processing using a single system-on-chip can simplify software integration and resource management but concentrates power density and may require aggressive thermal management. Distributed processing partitions the workload among microcontrollers or dedicated accelerators placed near sensors, reducing cable bandwidth and enabling localized preprocessing. However, distributed architectures require synchronization and robust communication links to ensure consistent state estimation. Under strict power constraints, both approaches may be operated at reduced voltage and frequency, trading throughput for efficiency.

Algorithm selection is tightly coupled to architecture. Lightweight model-based algorithms, including horizon detection, background modeling, and morphological filtering, can operate with low computational overhead but may be less robust across diverse backgrounds and lighting [25]. Convolutional neural networks trained on representative data sets can provide improved detection performance at the cost of significantly higher computation, memory, and implementation complexity. Pruning, quantization, and architectural search can reduce these costs, but aggressive compression may degrade performance in precisely those edge cases where detection is most critical. Feasibility considerations therefore involve quantifying whether an embedded platform with a given power envelope can sustain the computational load of a suitably robust algorithm at the frame rates and resolutions necessary for timely detection.

The final component, decision logic, maps detection and tracking outputs into avoidance actions. It must be implemented in a manner that is stable and predictable, as oscillatory or overly sensitive responses can be detrimental in dense airspace. Decision thresholds, temporal filters, and track confirmation logic must be tuned to balance false alarm rates and missed detections while respecting the limitations of the flight control system. This logic often interacts with other autonomy functions, such as mission planning and geofencing, and its complexity is not negligible [26]. Comprehensive feasibility analysis must regard detect-and-avoid as a system-level capability whose components jointly consume resources and whose interactions may expose emergent behavior.

Calibration and stability of multi-camera systems are particularly relevant for small unmanned aircraft. Relative orientations and positions of cameras must be known with sufficient accuracy that detected bearings can be fused into consistent three-dimensional hypotheses. Flexible airframes and mounting structures subject to vibration and temperature variations can cause slow drifts or rapid perturbations in calibration parameters.

Online calibration strategies can compensate for some of these effects but require additional computation and can be vulnerable to overfitting in low-feature environments. In contrast, rigid monocular installations minimize these complexities but reduce opportunities for geometric triangulation. The need for stable calibration thus exerts additional pressure toward mechanically simple configurations consistent with size and weight limits. [27]

Dynamic range and motion characteristics of the sensor frontend also affect detect-and-avoid performance. Rolling-shutter cameras, which dominate in compact low-power devices, can introduce geometric distortions when the platform or scene exhibits rapid motion, potentially degrading small-object localization. Global-shutter devices alleviate this issue but have historically incurred higher cost and power consumption. Similarly, exposure settings tuned for bright sky can saturate ground regions and obscure low-contrast targets, whereas settings favoring darker scenes may reduce sensitivity to distant aircraft against bright backgrounds. Adaptive exposure and high dynamic range techniques mitigate these effects at the cost of additional processing and potential temporal artifacts. Under strict resource budgets, design choices must balance these trade-offs to ensure that the effective signal presented to detection algorithms remains informative across expected operating conditions.

4 Mathematical Modeling of Detection, Tracking, and Collision Risk

A more detailed mathematical framework can formalize the relationship between vision system design and detect-and-avoid performance. Let the ownship and intruder states follow deterministic kinematics over short intervals, with relative position $r(t)$ and relative velocity $u(t)$ as before [28]. Under a constant relative velocity assumption, the time to closest approach τ and miss distance magnitude d can be expressed as

$$\tau = -\frac{r^\top u}{u^\top u},$$

$$d^2 = r^\top r - \frac{(r^\top u)^2}{u^\top u}.$$

These expressions provide a compact mapping from relative state to collision risk and underpin many detect-and-avoid logics. Vision-based sensing, however, typically observes only the bearing unit vector $b(t) = r(t)/\|r(t)\|$ projected into the camera frame, along with image intensity patterns that encode target size and appearance. Thus, the system must infer $r(t)$ and $u(t)$, or at least τ and d, from partial and noisy observations.

Assume that at discrete times t_k the system obtains measurements z_k comprising estimated bearing and possibly additional cues such as apparent scale. The observation model can be written abstractly as $z_k = h(x_k) + v_k$, with x_k denoting the relative state and v_k observation noise. A recursive estimator produces $\hat{x}_k$ with covariance P_k. For purposes of feasibility analysis, detailed estimator structure can be encapsulated by the resulting distributions of $\hat{\tau}_k$ and $\hat{d}_k$ given the true encounter. A conflict is declared when $\hat{\tau}_k$ lies within prescribed bounds and $\hat{d}_k$ falls below a threshold. Let E denote the event that a true collision would occur absent avoidance and A the event that an avoidance maneuver is initiated. The probability that the system prevents collision can be approximated as

$$P(A \mid E) = P(R_d \geq R_{req}),$$

where R_d is the detection range and R_{req} is the minimum range needed to execute an effective avoidance maneuver.

To connect detection range to system parameters, consider a simplified relation for the minimum detectable angular size θ_{min}, governed by optics and algorithm sensitivity. For a target of characteristic dimension S at range R, the apparent angular size is approximately S/R [29]. Detection requires $S/R \geq \theta_{min}$, implying $R \leq S/\theta_{min}$. The quantity θ_{min} depends on pixel pitch, lens focal length, and signal-to-noise ratio. For a focal length f and pixel pitch p, the angular sampling is approximately p/f, and a target should subtend a minimum number of pixels $n_{p,min}$. This yields $\theta_{min} \approx n_{p,min}p/f$ and thus

$$R_{max} \approx \frac{Sf}{n_{p,min}p}.$$

This expression, while idealized, illustrates how limited focal length and resolution reduce maximum detection range for small intruders, especially when size, weight, and power constraints restrict the use of large optics or high-resolution sensors.

False alarms and missed detections arise from background clutter and noise. If potential detections are generated as random events with rate λ_{FA} per solid angle and confirmed as tracks over a time window, the probability of at least one false track in a given interval scales with both λ_{FA} and the size of the observed region. Elevated false alarm rates may prompt conservative tuning that delays track confirmation, effectively increasing R_{req}. Conversely, aggressive tuning can reduce missed detections at the cost of more frequent false maneuvers. A stylized balance may be expressed as

$$P_{MD} \approx g(R_d; \theta_{min}, SNR),$$

$$P_{FA} \approx h(\lambda_{FA}; T_w),$$

where g and h are monotone functions reflecting algorithm behavior and T_w is the confirmation window. While exact forms depend on implementation, these relationships emphasize that hardware-imposed limits on resolution and sensitivity constrain achievable trade-offs between P_{MD} and P_{FA}.

The collision risk over an ensemble of encounters can be approximated by integrating residual collision probabilities conditioned on detect-and-avoid behavior. Let $f_E(r, u)$ describe the distribution of relative states for potential encounters. For each state, the associated detection range distribution and avoidance outcome define a conditional collision probability. The overall residual risk is then [30]

$$P_C = \int \psi(r, u)\, f_E(r, u)\, dr\, du,$$

where ψ encodes the outcome of detection, tracking, and avoidance given system characteristics. Feasibility in this sense involves demonstrating that for relevant f_E, realistic vision hardware and algorithms yield P_C below acceptable thresholds in the intended operational context, recognizing that such thresholds and distributions are externally specified rather than inherent to the technology.

A more refined treatment of detection performance introduces signal-to-noise ratio models for small targets in clutter. Let the average contrast between target and background in the relevant region of interest be C_t, the sensor noise standard deviation be σ_n, and the effective integration across N pixels contribute to detection. Then a simplified per-frame detection statistic may scale with

$$SNR \approx \frac{C_t\sqrt{N}}{\sigma_n}.$$

Assuming a threshold SNR_{min} for reliable detection, the minimum detectable apparent size and thus maximum detection range become functions of illumination, atmospheric attenuation, and sensor characteristics. In practice, C_t and σ_n vary with range and environment, so R_d is better treated as a random variable conditioned on these factors than as a fixed value. Monte Carlo or analytical approximations can then be employed to derive distributions F_{R_d} under specified conditions.

The mapping from detection to avoidance can be formalized through kinematic constraints. Consider an ownship capable of lateral acceleration up to a_{max} and operating at speed V_o. To achieve a miss distance not less than D_{min} against a head-on intruder at relative speed V_c, the approximate required warning time T_{req} for a coordinated turn maneuver can be bounded by

$$T_{req} \approx \sqrt{\frac{2D_{min}}{a_{max}}}.$$

The corresponding required detection range is then

$$R_{req} \approx V_c(T_p + T_{req}),$$

with T_p denoting combined sensing and processing latency. By comparing empirical or simulated R_d distributions with R_{req} across encounter classes, one obtains feasibility conditions expressed in terms of inequalities on tail probabilities such as $P(R_d \geq R_{req})$. This representation is convenient because it decouples detailed algorithm structure from higher-level performance assessment, provided the detection range statistics are measured or modeled credibly.

Uncertainties in intruder behavior can be incorporated by replacing deterministic closure rates with distributions and by modeling deviations from straight-line flight [31]. For example, if intruder trajectories experience bounded accelerations or execute turns, the relative motion model underlying τ and d must be adapted. State estimation frameworks can incorporate such uncertainties through process noise, leading to broader distributions for $\hat{\tau}$ and $\hat{d}$. The detect-and-avoid decision thresholds may then be selected to maintain desired probabilities of timely maneuver initiation across the uncertainty set, influencing both P_{MD} and P_{FA}. Again, size, weight, and power constraints influence how aggressively such estimation and decision logic can be implemented on embedded hardware.

5 Algorithmic Trade-offs Under Size, Weight, and Power Limitations

Algorithmic design for onboard vision-based detect-and-avoid must reconcile performance goals with the constraints of embedded implementation. Let C denote the sustained computational capacity of the processing hardware measured in operations per second allocated to detect-and-avoid. For a model that requires L arithmetic operations per frame at resolution and architecture of interest, operating at frame rate f demands $Lf \leq C$. Additional computations for tracking, data fusion, and control must also be accommodated within this budget. This inequality constrains the feasible combinations of model complexity, resolution, and frame rate. In low-power systems where C is modest, practical operation may necessitate reduced resolution or simpler models, which in turn influence detection range and robustness. [32]

Power consumption can be modeled in terms of activity and voltage. Let the processor consume dynamic power approximately proportional to an activity factor, effective capacitance, the square of supply voltage, and clock frequency. Lowering voltage and frequency reduces power but also reduces maximum C. Thermal constraints impose an upper bound on sustained power, which for small unmanned aircraft may be only a few watts. Given this, designers may operate near an efficiency point where additional computational margin is limited. Any algorithm that demands frequent peaks near the upper limit risks throttling or instability. Consequently, detect-and-avoid implementations often favor architectures with predictable and bounded computational load, such as fixed-depth networks and deterministic processing pipelines, over highly dynamic schemes whose resource demands can spike unpredictably. [33]

Latency emerges directly from these considerations. If processing of each frame incurs time T_p, then the total effective warning time available for avoidance is reduced from T_w to $T_w - T_p$. In high-speed encounters, even tens of milliseconds can be relevant. Suppose the minimum required warning time for an avoidance maneuver, given platform performance and desired miss distance D_{min}, is T_{req}. Then feasibility requires $R_d \geq V_c(T_p + T_{req})$ with high probability, where V_c is closure rate. For modest closure rates and agile platforms, this condition can be satisfied with moderate detection ranges. For higher closure rates or less agile platforms, margins become tight, and algorithmic or hardware adjustments are needed. Processing pipelines that enable early rejection of obvious non-threat regions can reduce effective L, thereby reducing T_p and improving feasibility.

Memory bandwidth and storage also play roles. High-resolution imagery at elevated frame rates requires substantial bandwidth for reading sensor data, writing intermediate results, and storing model parameters. Limited bandwidth can cause contention that increases latency or forces lower frame rates [34]. Compression techniques or region-of-interest processing can mitigate bandwidth demand but introduce additional complexity and potential degradation of small-object fidelity. Under size, weight, and power constraints, memory subsystems with wide buses and large caches may be impractical, necessitating tight coupling between algorithm design and memory architecture.

Robustness to implementation imperfections is another dimension of feasibility. Quantization of network weights and activations to low bit-widths can yield substantial energy savings but may affect detection sensitivity to subtle cues associated with distant or low-contrast intruders. Fixed-point implementations must be carefully scaled to avoid saturation and underflow across operating conditions. In addition, software complexity must be kept within limits that support verification, validation, and runtime monitoring appropriate for safety-related functions. While these considerations are not unique to vision-based detect-and-avoid, they are sharpened by the tight margins imposed by small unmanned aircraft platforms and by the need to maintain stable operation across heterogeneous mission profiles.

Real-time scheduling of vision workloads on embedded platforms introduces further considerations [35]. Pipelines often consist of sensor readout, image preprocessing, feature extraction, candidate generation, classification, and tracking updates, each with associated deadlines. Under tight power and computational budgets, designers may assign static time slots and prioritize tasks that are safety critical, such as conflict detection, over secondary tasks such as high-fidelity mapping. Preemption or dynamic frequency scaling in response to fluctuating load can stabilize performance but complices timing analysis. For detect-and-avoid, where worst-case latency bounds are central to safety arguments, predictable execution patterns are advantageous, even if they entail conservative resource allocation.

The choice between end-to-end learning-based pipelines and modular designs also has implications. End-to-end approaches can, in principle, optimize detection and classification jointly, but may obscure internal failure modes and pose challenges for incremental validation. Modular designs allow individual components to be assessed and tuned, but may forgo some potential performance gains from joint optimization [36]. Within small unmanned aircraft constraints, modular decompositions that combine lightweight motion-based pre-filters with focused application of heavier classifiers to selected regions appear compatible with both verification needs and resource limits. Such architectures can reduce average computational load while preserving responsiveness to small, distant targets, provided that pre-filters maintain high recall on relevant signatures.

Fault tolerance and runtime monitoring complete the algorithmic picture. Single points of failure in vision processing pipelines may expose operations to undetected degradation due to, for example, misconfigured exposure settings or corrupted model parameters. Lightweight consistency checks, watchdog timers, and self-test routines can be implemented to detect gross anomalies without incurring prohibitive overhead. Where resources permit, diversity in models or implementations can mitigate correlated failure modes, though full redundancy is typically unrealistic within small unmanned aircraft size and weight envelopes. These mechanisms do not eliminate uncertainty but contribute to bounding it, which is germane to feasibility assessments grounded in conservative assumptions. [37]

From a design methodology point of view, exploring algorithm configurations under parametric sweeps of available compute and power can reveal regimes in which diminishing returns set in. For example, increasing network depth beyond a certain point may yield marginal improvements in detection probability at the cost of disproportionate increases in latency and energy per frame, eroding overall feasibility. Conversely, carefully tuned shallow networks or hybrid schemes that exploit temporal consistency can approach similar performance with substantially lower resource demands. Embedding such analyses early in platform design can prevent late-stage integration challenges where detect-and-avoid requirements conflict with fixed hardware limitations.

6 Experimental Evaluation and Synthetic Campaign Design

Quantitative evaluation of vision-based detect-and-avoid systems under realistic size, weight, and power constraints requires experimentation strategies that probe both hardware and algorithm behavior. A combined approach often employs software-in-the-loop simulation, hardware-in-the-loop emulation, and flight testing. In software-in-the-loop configurations, representative encounter geometries and environmental conditions are rendered to produce synthetic imagery that spans variations in range, aspect, illumination, background, and target type. Vision algorithms execute on general-purpose hardware to estimate achievable detection performance and to explore parameter sensitivities [38]. While such evaluations do not capture all embedded constraints, they provide a controlled environment for initial screening of architectures and models.

Hardware-in-the-loop experiments map these evaluations closer to operational reality by running candidate algorithms on target or equivalent embedded processors while feeding them prerecorded or synthetic imagery. This allows direct measurement of processing latency, power consumption, and thermal behavior as functions of workload, as well as observation of any throttling or performance variation induced by resource limits. By systematically varying image resolution, frame rate, and algorithmic parameters, one can chart feasible operating regions in which throughput and latency requirements are met without violating power or thermal constraints. Mapping these regions onto detection performance derived from the imagery enables identification of configurations that are both computationally sustainable and operationally meaningful.

Flight testing remains essential for exposing systems to unmodeled effects, including platform vibrations, dynamic lighting, atmospheric scattering, and unexpected background structures. Small unmanned aircraft equipped with candidate vision systems can be flown in instrumented environments with cooperative intruders following scripted trajectories, providing ground-truth data for post hoc analysis [39]. Safety considerations require conservative separation buffers and independent monitoring, limiting the range of conditions that can be explored directly. Nonetheless, flight tests can validate key aspects of the sensing chain, such as stability of calibration, resilience to motion blur, and sensitivity to practical integration issues like lens contamination and cable flexure. These insights feed back into both algorithm design and synthetic modeling.

Evaluation metrics in such campaigns typically include detection probability as a function of range, false alarm rate per unit time or solid angle, track continuity, and latency from image acquisition to conflict declaration. By aggregating results across encounter classes and environmental conditions, one can estimate empirical detection range distributions and associated decision latencies. These quantities, combined with models of platform maneuverability and encounter kinematics, allow construction of approximate residual risk estimates analogous to the earlier integral expression for collision probability. Importantly, these estimates are conditional on the particular hardware, algorithms, and operating assumptions tested [40]. Feasibility conclusions should therefore be framed in terms of the specific configurations and envelopes studied, rather than generalized without qualification to all small unmanned aircraft or all operations.

In constructing synthetic data sets, care is required to avoid optimistic biases. If training and evaluation data are drawn from similar distributions with limited variability, estimated detection performance may appear strong yet fail to generalize to novel environments. Including challenging cases such as cluttered urban skylines, sun glint, partial occlusions, and diverse aircraft textures is therefore important. Similarly, scenario sampling should avoid overrepresentation of benign geometries with long warning times, as this can mask weaknesses in rare but critical head-on or crossing encounters with limited time margins. From a feasibility perspective, emphasizing coverage of adverse and edge cases yields a more informative characterization, even if it reduces nominal performance metrics.

Measurements from hardware-in-the-loop and flight tests can be used to calibrate simulation models [41]. For example, observed relationships between ambient temperature, processor utilization, and throttling behavior inform power and thermal models, which then constrain allowable workloads in broader simulations. Likewise, empirical distributions of detection range and latency under measured visibility conditions can validate or refine analytic models. This iterative process aligns modeled feasibility assessments with observed system behavior, reducing the risk that conclusions rest on unrealistic assumptions.

A complementary dimension of evaluation involves the interaction between detect-and-avoid outputs and flight control. Avoidance maneuvers commanded by the vision system must integrate coherently with existing stabilization and navigation loops. Experiments can examine whether commanded maneuvers are faithfully executed given actuation limits, delays, and concurrent commands from higher-level mission planners. Situations in which aggressive avoidance demands conflict with vehicle stability or mission constraints are particularly salient, as they may reveal implicit assumptions about available maneuver authority. Understanding these interactions is integral to determining whether a vision-based system can, in practice, translate detections into effective risk mitigation on a given platform. [42]

7 Integrated Design Space and Feasibility Assessment

Bringing together system constraints, sensing architectures, mathematical modeling, and evaluation results, feasibility can be discussed in terms of an integrated design space spanned by key parameters. These include optical aperture and focal length, sensor resolution and frame rate, processor throughput and power consumption, platform maneuver capability, and characteristics of the operational environment. Each point in this space corresponds to a potential detect-and-avoid configuration whose performance can be approximated through the models and methods outlined earlier. The structure of this space is shaped by physical and practical limits: optics cannot exceed certain diameters without violating mass or drag constraints, processors cannot consume more than a bounded fraction of available electrical power, and platforms cannot consistently deliver accelerations beyond their mechanical capabilities.

Within this space, regions can be identified where detect-and-avoid performance is materially constrained by specific bottlenecks. For instance, combinations of low-resolution cameras and limited focal length may cap maximum detection range below values compatible with high closure rate encounters, regardless of algorithmic sophistication. Conversely, powerful processors operating at higher power levels might support complex algorithms and high-resolution input, but only at the expense of reduced endurance or increased thermal stress that may be unacceptable for some missions [43]. More favorable regions arise when modestly capable sensors and processors are mounted on agile platforms operating in structured environments with bounded closure rates and predictable backgrounds, enabling sufficient warning times without extreme hardware demands.

Trade-offs between different resource allocations can also be characterized. Increasing the number of cameras improves angular coverage and reduces blind zones, but adds mass and elevates processing load. Increasing focal length extends detection range for a given target size but narrows field of view, potentially increasing the probability that intruders approach from outside coverage arcs. Raising frame rate improves temporal resolution for tracking and reduces motion blur, but scales linearly with computational requirements. Under fixed size, weight, and power budgets, improvements along one dimension typically require reductions along another, and not all directions lead to net benefit. Feasibility assessment thus benefits from explicit multi-dimensional optimization frameworks that balance these competing effects with respect to representative encounter models and risk metrics. [44]

In such optimization, the role of operational constraints becomes explicit. If maximum relative speeds in a particular corridor are limited by procedure, the necessary detection range can be reduced, broadening the set of feasible onboard vision configurations. If operations are confined to daylight and good visibility, algorithmic and optical demands are relaxed relative to all-weather requirements. Conversely, if operations demand coexistence with faster crewed aircraft in uncontrolled environments, or if night operations are required, feasibility may hinge on more capable sensors or on augmenting vision with cooperative surveillance or alternative sensing modalities. Rather than assuming uniform applicability, a neutral assessment treats onboard vision-based detect-and-avoid as one design variable among several, whose appropriateness depends on context.

Illustrative case studies help make the integrated design space concrete. Consider a multirotor platform with maximum take-off mass near 5 kg, capable of lateral accelerations approximately 4 m/s^2 and typical cruise speeds near 12 m/s. Suppose the payload budget permits allocation of 150 g and 8 W to a detect-and-avoid subsystem, supporting two lightweight global-shutter cameras and an embedded processor delivering effective capacity in the low gigaflops range. Encounter modeling focused on other small unmanned aircraft and light helicopters with closure rates below 50 m/s may indicate that detection ranges on the order of several hundred meters are sufficient to enable avoidance [45]. Synthetic and flight experiments could then evaluate whether candidate algorithms on the selected hardware achieve such ranges with acceptable false alarm rates across relevant conditions. If so, this configuration can be regarded as within a feasible region for vision-based detect-and-avoid.

In contrast, consider a small fixed-wing delivery aircraft of similar mass operating at 30 m/s, sharing uncontrolled airspace with crewed aircraft at significantly higher closure rates. For such a platform, required detection ranges to support lateral or vertical maneuvers can extend into the kilometer regime. Under tight size, weight, and power limits, achievable optical apertures and sensor resolutions may yield reliable detection only at substantially shorter distances. Unless complemented by cooperative surveillance or other sensing modalities, onboard vision alone may then provide limited margin. These examples illustrate how feasibility is contingent on aligning system design with the specific kinematic and environmental envelope rather than presuming uniform applicability. [46]

As technologies evolve, the boundaries of feasible regions can shift. Improvements in sensor quantum efficiency, on-chip processing, and low-power accelerators can extend detection ranges or reduce latencies without proportionally increasing power consumption. Algorithmic advances in efficient network architectures and training methodologies can enhance robustness at fixed computational budgets. However, variability in real-world environments and traffic behaviors persists, and structural limitations such as line-of-sight dependence remain. Consequently, updated feasibility analyses should continue to adopt conservative, data-informed assumptions and should present results in terms of conditional statements tied to explicit configurations and operational profiles.

Finally, it is useful to note that many of the dependencies in this design space are smooth rather than abrupt. Small incremental increases in available power, modest improvements in optical quality, or incremental gains in algorithmic efficiency can collectively shift a configuration from marginal to acceptable performance [47]. Similarly, relatively small degradations in visibility, calibration accuracy, or processor health can erode margins. Feasibility evaluations that acknowledge these gradients can inform design choices that favor graceful degradation and resilience, such as selecting operating points with surplus computational capacity, designing optics with moderate reserve in resolution, or adopting control policies that maintain maneuver authority explicitly reserved for avoidance. In this way, onboard vision-based detect-and-avoid can be incorporated as part of a broader engineering approach that prioritizes transparent margins and incremental improvement over singular dependencies on any individual component or performance claim.

Another dimension of the design space concerns integration with higher-level traffic management concepts. If small unmanned aircraft operations are constrained to designated corridors or altitudes where background clutter and traffic patterns are relatively structured, vision-based detect-and-avoid algorithms can exploit these regularities, for example by tailoring detection thresholds to dominant approach angles or by allocating higher resolution to sectors with elevated conflict likelihood. Conversely, if operations occur in fully unstructured environments with broad distributions of intruder approach directions and behaviors, conservative tuning may be necessary, potentially increasing false alarm rates and computational burden. The feasibility of onboard vision is therefore linked not only to physical hardware parameters but also to the extent to which operational frameworks can reduce uncertainty in the encounter space. Reasoned combinations of modest procedural constraints with resource-aware vision system designs can, in many cases, produce more favorable feasibility outcomes than attempts to rely on either element in isolation. [48]

8 Conclusion

The feasibility of onboard vision-based detect-and-avoid for small unmanned aircraft systems under stringent size, weight, and power constraints arises from the interdependent balance between sensing architecture, computational design, dynamic performance of the airframe, and the operational context within which the system must function. Vision-based sensing offers a compact and passive means to achieve angular awareness of surrounding traffic without reliance on cooperative transponders or external infrastructure. The potential of such systems lies in their ability to detect and track non-cooperative objects using monocular or stereo imagery, translating visual information into actionable collision avoidance decisions. However, the practical effectiveness of these systems is not intrinsic to the technology itself; it is defined by the achievable detection ranges, the latency with which data can be processed into decisions, the reliability of classification under variable illumination, and the false alarm behavior under complex backgrounds. Environmental variability, encompassing factors such as haze, glare, and clutter, often imposes stochastic limitations on detection probabilities that cannot be completely eliminated by algorithmic sophistication alone.

Mathematical models connecting detection range distributions, encounter kinematics, and avoidance maneuver dynamics allow the estimation of whether the temporal margins available for avoidance are sufficient under different encounter geometries. In simplified form, the expected available warning time can be represented as the ratio of the detection range to the closure rate between aircraft, which must exceed the sum of sensing and control latencies plus the time required to execute a maneuver of given acceleration limits [49]. If R_d denotes the detection range and V_c the relative closure velocity, the warning time T_w satisfies $T_w = R_d/V_c$. Avoidance feasibility requires that T_w exceed a threshold composed of processing latency T_p and maneuver time T_m, implying $R_d > V_c(T_p+T_m)$. For small unmanned aircraft operating at moderate closure rates, achievable detection ranges of several hundred meters can yield feasible warning times when computational and control latencies are minimized. Yet for faster encounters or degraded environmental conditions, margins collapse quickly. This illustrates how physical parameters, algorithmic efficiency, and platform agility jointly determine whether onboard vision sensing can maintain safe separation.

The algorithmic dimension introduces a dense network of trade-offs. Increasing model complexity can enhance detection robustness across diverse backgrounds but scales computation and power requirements nonlinearly. Input resolution directly affects detection range because smaller targets subtend fewer pixels; however, higher resolutions increase memory traffic and computational load. Similarly, increasing frame rate reduces temporal aliasing and improves responsiveness, but also raises power consumption and heat generation [50]. Quantization precision in learned models influences inference accuracy and energy efficiency, with low-bit quantization reducing power draw but potentially degrading sensitivity to small or low-contrast targets. These interdependencies define a narrow region of feasible configurations where accuracy, latency, and energy coexist in acceptable balance. Within this region, co-design of hardware and software is imperative. Algorithmic architectures must be tailored to the capabilities of specific embedded processors, exploiting parallelism and minimizing data movement to maintain real-time performance without exceeding thermal limits.

Embedded processing platforms impose strict upper bounds on throughput and deterministic timing behavior. Small unmanned aircraft often rely on processors delivering only a few gigaflops of sustained performance at power levels below 10 W. To meet these limits, detect-and-avoid algorithms must employ efficiency strategies such as lightweight convolutional architectures, selective region-of-interest processing, or temporal filtering that leverages prior frame information to reduce redundant computation [51]. The processing chain must be designed for bounded latency; jitter in inference time can produce unpredictable decision delays that degrade safety margins. Real-time operation requires that each stageimage acquisition, preprocessing, feature extraction, detection, tracking, and decision logicexecute within a predictable time window synchronized with the aircrafts control loops.

Synthetic imagery and high-fidelity simulation provide a first layer of evaluation, enabling systematic variation of conditions that would be difficult to reproduce in flight. Synthetic campaigns can explore relationships between apparent target size, contrast, and detection probability, generating empirical detection range distributions. Hardware-in-the-loop experiments bridge the gap between simulation and flight by running these pipelines on actual embedded hardware under controlled workloads, quantifying latencies and thermal effects. These experiments can reveal subtle interactions such as frequency throttling under sustained load, which might not appear in simulation but strongly influence real-time feasibility. Targeted flight testing then validates the full sensing and control chain in realistic conditions, exposing effects of vibration, lighting transitions, and attitude-induced camera motion. Together, these methodologies constitute a layered evaluation approach that quantifies detection probability, false alarm rates, latency, and overall residual collision risk across representative operational envelopes. [52]

Residual collision risk serves as the ultimate metric connecting detection performance to operational safety. The risk can be formalized as the probability that an encounter results in a loss of separation given the sensing and avoidance system in place. This probability depends on both physical and decision-level factors. Let P_{MD} denote the missed detection probability, P_{FA} the false alarm probability, and $P_{A|D}$ the conditional probability of successful avoidance given a detection. Then the residual collision probability P_C over an ensemble of encounters may be approximated as $P_C \approx (1 - P_{A|D})(1 - P_{FA})P_{MD}$. Although simplistic, this representation highlights that both sensing and control components must perform reliably for risk to remain acceptably low. Vision-based systems reduce P_{MD} through improved algorithms and optics but may increase P_{FA} under clutter. Hardware and software co-optimization seeks to minimize the composite product of these probabilities under fixed resource budgets.

Evaluating such metrics requires consistent definitions of operational envelopes. For example, a small quadrotor flying at 12 m/s encounters relative closure rates of 24 m/s in symmetric head-on engagements with similar aircraft [53]. If the systems effective detection range is 250 m and total response latency 1.5 s, the available warning time is approximately 10.4 s, leaving roughly 8.9 s for avoidance after latency deductionan adequate margin for most lateral or vertical maneuvers given typical acceleration capacities near 4 m/s^2. However, under degraded visibility reducing detection range to 100 m, warning time falls below 4.2 s, and after accounting for latency, maneuvering time drops below 3 s, which is marginal. These numeric illustrations underscore the sensitivity of feasibility to environmental and system parameters. Even modest reductions in range or increases in latency can produce disproportionate decreases in avoidance margin.

False alarm management remains a critical aspect of operational feasibility. Excessive false alarms can cause unnecessary maneuvers, reducing mission efficiency and potentially introducing new risks through erratic flight behavior. Practical systems must balance conservatism in detection with operational stability. Temporal confirmation logic, which requires consistent detection across multiple frames before declaring a threat, reduces false alarms but increases detection latency [54]. The trade-off between false alarm suppression and timely response must therefore be carefully tuned based on encounter statistics and mission priorities. Analytical modeling and simulation help quantify these trade-offs under different parameterizations, guiding system-level design.

Environmental robustness also determines whether onboard vision-based systems can operate across diverse contexts. Variations in ambient brightness, dynamic range, and background texture significantly affect detection reliability. Adaptive exposure control, dynamic thresholding, and contrast normalization can partially compensate but consume computational resources and may not fully eliminate sensitivity to extreme conditions. Cloud shadows, glare from low-angle sunlight, or reflective surfaces can produce false positives. Sensor placement and optical coatings help mitigate some of these effects but cannot guarantee invariance. Consequently, realistic feasibility assessments must account for degraded conditions and establish operational envelopesranges of illumination and weather within which specified performance can be maintained with high probability. [55]

Integration with flight control systems introduces further complexity. Avoidance maneuvers derived from visual detection must be compatible with vehicle dynamics and mission objectives. Multirotors can execute rapid lateral translations, while fixed-wing aircraft require banked turns and experience greater inertia. The control interface must therefore interpret avoidance commands within the vehicles flight envelope, ensuring stability and preventing excessive control deflections. Timing synchronization between perception and control loops is essential; even minor mismatches can lead to phase lags that reduce maneuver effectiveness. Testing under coupled perceptioncontrol operation verifies that detection events propagate through the control chain without excessive latency or oscillation.

The integrated assessment of all these factors reveals that vision-based detect-and-avoid is neither universally sufficient nor categorically infeasible [56]. It occupies a middle ground where feasibility depends on explicit matching of components and context. For platforms with moderate speeds, agile maneuvering, and operations constrained to daylight and clear weather, onboard vision sensing can deliver meaningful reductions in collision risk. For faster aircraft or operations requiring all-weather performance, additional sensing modalities or cooperative aids become necessary. The value of vision lies not in replacing all other systems but in augmenting the situational awareness of small platforms operating autonomously or semi-autonomously in complex environments.

Continued advances in image sensors, optics, and embedded processing promise gradual expansion of the feasible design space. Improved quantum efficiency and on-chip integration reduce power consumption while increasing sensitivity. Dedicated neural accelerators and low-power graphics processors enable real-time inference with modest energy budgets [57]. Algorithmic innovationssuch as event-based processing, temporal integration, and efficient model compressionextend operational viability without violating size, weight, and power constraints. Nevertheless, each increment in performance must be validated through systematic evaluation rather than assumed from simulation or laboratory data. Robust methodologies encompassing synthetic, hardware-in-the-loop, and flight domains remain the foundation for understanding both strengths and limitations.

Ultimately, the feasibility of onboard vision-based detect-and-avoid rests on pragmatic engineering alignment. Sensors must provide sufficient optical fidelity, processors must deliver predictable latency, algorithms must be computationally tractable, and the platform must sustain necessary maneuver authority. Each element interacts multiplicatively rather than additively; a deficiency in any component propagates through the system, constraining the whole. Recognizing this interdependence leads to design strategies emphasizing co-optimization and balanced margins rather than singular focus on any one factor. With disciplined integration and empirically grounded evaluation, onboard vision sensing can serve as a viable component of layered separation assurance architectures for small unmanned aircraft, contributing to safe coexistence in shared low-altitude airspace while acknowledging the intrinsic limits imposed by physics and available resources [58].

References

[1] S. Gotovac, D. Zelenika,. Marui, and D. Boi-tuli, "Visual-based person detection for search-and-rescue with uas: Humans vs. machine learning algorithm," *Remote Sensing*, vol. 12, no. 20, pp. 3295–, Oct. 10, 2020.

[2] I. Daugla, J. S. Visockiene, and J. Kumpiene, "Detection and analysis of methane emissions from a landfill using unmanned aerial drone systems and semiconductor sensors," *Detritus*, vol. 10, no. 10, pp. 127–138, May 8, 2020.

[3] M. R. James et al., "Volcanological applications of unoccupied aircraft systems (uas): Developments, strategies, and future challenges," *Volcanica*, vol. 3, no. 1, pp. 67–114, Apr. 9, 2020.

[4] G. A. Davis, M. L. Giles, and D. M. Tate, "Complexity in an unexpected place: Quantities in selected acquisition reports," *Defense Acquisition Research Journal*, vol. 27, no. 91, pp. 28–59, Jan. 1, 2020.

[5] J. H. F. Castenschiold, T. Bregnballe, D. Bruhn, and C. Pertoldi, "Unmanned aircraft systems as a powerful tool to detect fine-scale spatial positioning and interactions between waterbirds at high-tide roosts.," *Animals : an open access journal from MDPI*, vol. 12, no. 8, pp. 947–947, Apr. 7, 2022.

[6] A. Westergaard-Nielsen, T. Balstrøm, U. A. Treier, S. Normand, and B. Elberling, "Estimating meltwater retention and associated nitrate redistribution during snowmelt in an arctic tundra landscape," *Environmental Research Letters*, vol. 15, no. 3, pp. 034 025–, Feb. 21, 2020.

[7] C. Huang, H. Zhang, and J. Zhao, "High-efficiency determination of coastline by combination of tidal level and coastal zone dem from uav tilt photogrammetry," *Remote Sensing*, vol. 12, no. 14, pp. 2189–, Jul. 8, 2020.

[8] D. Jusufovski, C. Saavedra, and A. Kuparinen, "Competition between marine mammals and fisheries in contemporary harvested marine ecosystems," *Marine Ecology Progress Series*, vol. 627, pp. 207–232, Sep. 26, 2019.

[9] Z. Song, H. Li, C. Chen, X. Zhou, and F. Xu, "Coordinated standoff tracking of moving targets using differential geometry," *Journal of Zhejiang University SCIENCE C*, vol. 15, no. 4, pp. 284–292, Apr. 12, 2014.

[10] C. Patterson, W. Koski, P. Pace, B. McLuckie, and D. M. Bird, "Evaluation of an unmanned aircraft system for detecting surrogate caribou targets in labrador1," *Journal of Unmanned Vehicle Systems*, vol. 4, no. 1, pp. 53–69, Mar. 1, 2016.

[11] J. Linchant et al., “Uas imagery reveals new survey opportunities for counting hippos,” *PloS one*, vol. 13, no. 11, e0206413–, Nov. 14, 2018.

[12] J. Schumann et al., “Software health management with bayesian networks,” *Innovations in Systems and Software Engineering*, vol. 9, no. 4, pp. 271–292, Jun. 21, 2013.

[13] C. Hoffmann, C. Weise, T. Koch, and K. Pauly, “From uas data acquisition to actionable information how an end-to-end solution helps oil palm plantation operators to perform a more sustainable plantation management,” *The International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, vol. XLI-B1, pp. 1113–1120, Jun. 6, 2016.

[14] J. J. Ely et al., “Establishing a disruptive new capability for nasa to fly uav’s into hazardous conditions,” *SPIE Proceedings*, vol. 9468, pp. 198–209, May 22, 2015.

[15] J. Lundberg, M. Arvola, and K. L. Palmerius, “Human autonomy in future drone traffic: Joint human-ai control in temporal cognitive work.,” *Frontiers in artificial intelligence*, vol. 4, pp. 704 082–, Jul. 20, 2021.

[16] U. Press, “Parameters summer 2021,” *The US Army War College Quarterly: Parameters*, vol. 51, no. 2, May 18, 2021.

[17] A. Shelley, “A framework for counter-unmanned aircraft system regulation in new zealand,” *Policy Quarterly*, vol. 14, no. 3, Aug. 13, 2018.

[18] D. Zhang et al., “Toward implementation of combined incompatible and sterile insect techniques for mosquito control: Optimized chilling conditions for handling aedes albopictus male adults prior to release.,” *PLoS neglected tropical diseases*, vol. 14, no. 9, e0008561–, Sep. 3, 2020.

[19] D. Brus, J. Gustafsson, V. Vakkari, O. Kemppinen, G. de Boer, and A. Hirsikko, “Measurement report: Properties of aerosol and gases in the vertical profile during the lapse-rate campaign,” *Atmospheric Chemistry and Physics*, vol. 21, no. 1, pp. 517–533, Jan. 15, 2021.

[20] R. Xu, C. Li, and A. H. Paterson, “Multispectral imaging and unmanned aerial systems for cotton plant phenotyping.,” *PloS one*, vol. 14, no. 2, e0205083–, Feb. 27, 2019.

[21] A. C. Canolla, M. B. Jamoom, and B. Pervan, “Unmanned aircraft systems detect and avoid sensor hybrid estimation error analysis,” in *17th AIAA Aviation Technology, Integration, and Operations Conference*, 2017, p. 4384.

[22] A. Sigala and B. T. Langhals, “Applications of unmanned aerial systems (uas): A delphi study projecting future uas missions and relevant challenges,” *Drones*, vol. 4, no. 1, pp. 8–, Mar. 10, 2020.

[23] U. Papa and S. Ponte, “Preliminary design of an unmanned aircraft system for aircraft general visual inspection,” *Electronics*, vol. 7, no. 12, pp. 435–, Dec. 14, 2018.

[24] M. Koeva, M. Muneza, C. Gevaert, M. Gerke, and F. Nex, “Using uavs for map creation and updating : A case study in rwanda,” *Survey Review*, vol. 50, no. 361, pp. 312–325, Dec. 30, 2016.

[25] M. B. Jamoom, A. Canolla, B. Pervan, and M. Joerger, “Unmanned aircraft system sense and avoid integrity: Intruder linear accelerations and analysis,” *Journal of Aerospace Information Systems*, vol. 14, no. 1, pp. 53–67, 2017.

[26] B. S. Acharya et al., "Unmanned aerial vehicles in hydrology and water management: Applications, challenges, and perspectives," *Water Resources Research*, vol. 57, no. 11, Nov. 18, 2021.

[27] G. T. Raber and S. R. Schill, "Reef rover: A low-cost small autonomous unmanned surface vehicle (usv) for mapping and monitoring coral reefs," *Drones*, vol. 3, no. 2, pp. 38–, Apr. 17, 2019.

[28] A. Al-Dhaqm, R. A. Ikuesan, V. R. Kebande, S. A. Razak, and F. Ghabban, "Research challenges and opportunities in drone forensics models," *Electronics*, vol. 10, no. 13, pp. 1519–, Jun. 23, 2021.

[29] S. Zhao, F. Lin, K. Peng, X. Dong, B. M. Chen, and T. H. Lee, "Vision-aided estimation of attitude, velocity, and inertial measurement bias for uav stabilization," *Journal of Intelligent & Robotic Systems*, vol. 81, no. 3, pp. 531–549, Feb. 11, 2015.

[30] D. Frommholz, M. Linkiewicz, and A.-M. Poznanska, "Inlining 3d reconstruction, multi-source texture mapping and semantic analysis using oblique aerial imagery," *The International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, vol. XLI-B3, pp. 605–612, Jun. 9, 2016.

[31] B. Wang et al., "Detection and quantification of forest-agriculture ecotones caused by returning farmland to forest program using unmanned aircraft imagery," *Diversity*, vol. 14, no. 5, pp. 406–406, May 20, 2022.

[32] S. H. Haji et al., "Comparison of software defined networking with traditional networking," *Asian Journal of Research in Computer Science*, pp. 1–18, May 27, 2021.

[33] G. Ariante, S. Ponte, U. Papa, A. Greco, and G. D. Core, "Ground control system for uas safe landing area determination (slad) in urban air mobility operations.," *Sensors (Basel, Switzerland)*, vol. 22, no. 9, pp. 3226–3226, Apr. 22, 2022.

[34] B. Aydin, E. Selvi, J. Tao, and M. J. Starek, "Use of fire-extinguishing balls for a conceptual system of drone-assisted wildfire fighting," *Drones*, vol. 3, no. 1, pp. 17–, Feb. 12, 2019.

[35] A. Tabassum and W. H. Semke, "Uat ads-b data anomalies and the effect of flight parameters on dropout occurrences," *Data*, vol. 3, no. 2, pp. 19–, Jun. 8, 2018.

[36] V. Bhandari, "Use of technology in disaster management," *Unity Journal*, vol. 3, no. 1, pp. 292–304, Mar. 6, 2022.

[37] J. A. Besada et al., "Drone mission definition and implementation for automated infrastructure inspection using airborne sensors.," *Sensors (Basel, Switzerland)*, vol. 18, no. 4, pp. 1170–, Apr. 11, 2018.

[38] S. M. Hashemi, R. M. Botez, and G. Ghazi, "Blockchain pos and pow consensus algorithms for airspace management application to the uas-s4 ehécatl," *Algorithms*, vol. 16, no. 10, pp. 472–472, Oct. 7, 2023.

[39] Q. He et al., "Intelligent construction for the transportation infrastructure: A review," *Intelligent Transportation Infrastructure*, vol. 1, Sep. 1, 2022.

[40] H. Sun, H. Yan, M. Hassanalian, J. Zhang, and A. Abdelkefi, "Uav platforms for data acquisition and intervention practices in forestry: Towards more intelligent applications," *Aerospace*, vol. 10, no. 3, pp. 317–317, Mar. 22, 2023.

[41] R. Brownlow et al., "Methane mole fraction and 13c above and below the trade wind inversion at ascension island in air sampled by aerial robotics," *Geophysical Research Letters*, vol. 43, no. 22, Nov. 19, 2016.

[42] S. Düsing, B. Wehner, T. Müller, A. Stöcker, and A. Wiedensohler, “The effect of rapid relative humidity changes on fast filter-based aerosol-particle light-absorption measurements: Uncertainties and correction schemes,” *Atmospheric Measurement Techniques*, vol. 12, no. 11, pp. 5879–5895, Nov. 8, 2019.

[43] J. Alvarenga, N. I. Vitzilaios, K. P. Valavanis, and M. J. Rutherford, “Survey of unmanned helicopter model-based navigation and control techniques,” *Journal of Intelligent & Robotic Systems*, vol. 80, no. 1, pp. 87–138, Dec. 3, 2014.

[44] F. R. Fabresse, F. Caballero, I. Maza, and A. Ollero, “Robust range-only slam for unmanned aerial systems,” *Journal of Intelligent & Robotic Systems*, vol. 84, no. 1, pp. 297–310, Dec. 23, 2015.

[45] R. V. Petrescu, “Some aspects related to drones,” *American Journal of Engineering and Applied Sciences*, vol. 13, no. 3, pp. 531–562, Mar. 1, 2020.

[46] G. Verhoeven and C. Sevara, “Trying to break new ground in aerial archaeology,” *Remote Sensing*, vol. 8, no. 11, pp. 918–, Nov. 4, 2016.

[47] I. I. David, J. Dalina, R. D. S. Manzon, R. J. L. D. Lara, and M. J. M. Villar, “Portable 5.8ghz fpv system with smartaudio,” *Engineering and Technology Journal*, vol. 06, no. 05, pp. 884–888, May 15, 2021.

[48] K. Jeong, J. Kwon, S. L. Do, D. Lee, and S. Kim, “A synthetic review of uas-based facility condition monitoring,” *Drones*, vol. 6, no. 12, pp. 420–420, Dec. 15, 2022.

[49] N. Mohamad, A. Ahmad, M. F. A. Khanan, and A. H. M. Din, “Surface elevation changes estimation underneath mangrove canopy using snerl filtering algorithm and dod technique on uav-derived dsm data,” *ISPRS International Journal of Geo-Information*, vol. 11, no. 1, pp. 32–32, Dec. 31, 2021.

[50] F. Bonfante, M. D. L. D. Vedova, and P. Maggiore, “Evaluations on adequacy and utility of failure modes and effects and criticality analysis and fault tree analysis methodologies applied to civil rpas systems,” *MATEC Web of Conferences*, vol. 233, pp. 00 002–, Nov. 21, 2018.

[51] B. Mamaghani and C. Salvaggio, “Multispectral sensor calibration and characterization for suas remote sensing.,” *Sensors (Basel, Switzerland)*, vol. 19, no. 20, pp. 4453–, Oct. 14, 2019.

[52] V. V. Koryanov et al., “Concept development of control system for perspective unmanned aerial vehicles,” *MATEC Web of Conferences*, vol. 151, pp. 04 010–, Feb. 21, 2018.

[53] M. Gharibi, Z. Gharibi, R. Boutaba, and S. L. Waslander, “A density-based and lane-free microscopic traffic flow model applied to unmanned aerial vehicles,” *Drones*, vol. 5, no. 4, pp. 116–, Oct. 12, 2021.

[54] D. Safadinho, J. Ramos, R. Ribeiro, V. Filipe, J. Barroso, and A. Pereira, “Uav landing using computer vision techniques for human detection,” *Sensors (Basel, Switzerland)*, vol. 20, no. 3, pp. 613–, Jan. 22, 2020.

[55] S. Vadlamudi, “How artificial intelligence improves agricultural productivity and sustainability: A global thematic analysis,” *Asia Pacific Journal of Energy and Environment*, vol. 6, no. 2, pp. 91–100, Dec. 31, 2019.

[56] P. Tymków, G. Jóków, A. Walicka, M. Karpina, and A. Borkowski, “Identification of water body extent based on remote sensing data collected with unmanned aerial vehicle,” *Water*, vol. 11, no. 2, pp. 338–, Feb. 16, 2019.

[57] A. Straubinger, J. Michelmann, and T. Biehle, “Business model options for passenger urban air mobility,” *CEAS aeronautical journal*, vol. 12, no. 2, pp. 1–20, Apr. 9, 2021.

[58] T. Elderini, N. Kaabouch, and J. Neubert, “Space occupancy representation based on a bayesian model for unmanned aerial vehicles,” *Journal of Intelligent & Robotic Systems*, vol. 97, no. 2, pp. 399–410, May 28, 2019.