Toward Aerial Edge Intelligence: A Vision for Next-Generation Drones

Christopher Stewart, *The Ohio State University*
Kevin Butler, *University of Florida*
Sabine Brunswicker, *Purdue University*
Nils Aschenbruck, *Osnabrück University*
Alina Gerall, *Computing Research Association*
Srini Parthasarathy, *The Ohio State University*
Weisong Shi, *University of Delaware*
Deborah Silver, *Rutgers University*
Andrew Thurling, *Thurling Aero Consulting*
Ufuk Topcu, *University of Texas at Austin*

Abstract—Drone technology has advanced well beyond simple uncrewed flight; modern drones are capable compute engines with reliable wireless connectivity. At the same time, artificial intelligence that reasons and plans can now run on resource-lean devices, allowing drones and nearby edge infrastructure to manage computation locally or offload to data centers, whichever best advances mission objectives. These trends enable aerial edge intelligence*, an emerging capability that could enable autonomous fleets of uncrewed aircraft to execute complicated, multi-faceted missions in dynamic conditions. However, realizing aerial edge intelligence requires new computing capabilities for safe, scalable aerial operations. In December 2025, we convened a workshop of 47 experts. This article identifies properties of aerial edge intelligence that existing frameworks miss, shows how they create three critical challenges, and proposes a research agenda for computing on the fly.*

Index Terms: *Aerial edge intelligence, autonomous drone fleets, edge computing, UAV, fleet-level behavior, trust and security, computing on the fly.*

1. Introduction

Two concurrent technological trends are converging to create an emerging grand challenge in computing. The first is the maturation of drone platforms. Power demands have decreased dramatically [1], lightweight materials have extended flight times [2], and flight-control software has steadily improved [3], [4]. Uncrewed flight, once a grand challenge in its own right, is now a baseline standard. Modern drones are increasingly capable compute engines equipped with sensors, actuators, and reliable wireless connectivity and their accessibility is growing. U.S. regulators project that the commercial drone fleet will roughly triple within a decade, with beyond-visual-line-of-sight (BVLOS) operations becoming commonplace [5], [6]. The Department of Defense's Replicator initiative calls for manufacturing thousands of autonomous aerial systems in the coming years [7]. The second trend is the arrival of artificial intelligence that can reason and plan on resource-lean edge computing devices. Rapid advances in model pruning, compression, and distillation [8]–[10], and hardware-software co-design [11], [12] have made sophisticated inference and planning practical on the constrained resources available at far-

XXXX-XXX © IEEE
Digital Object Identifier 10.1109/XXX.0000.0000000

edge sites.

Together, these trends enable *aerial edge intelligence*, i.e., networked systems that control and coordinate multiple drones, sensors at edge sites, and compute resources across the cloud-to-edge continuum to execute complex missions in dynamic environments at runtime (colloquially, we refer to this capability as computing on the fly). Aerial edge intelligence represents a qualitative leap in drone capability, enabling commercial applications in agriculture, cargo, emergency response, ecology, and infrastructure inspection that are already moving from research prototypes to reliable services. Yet translating this potential into safe, scalable aerial operations is challenging. Existing edge and distributed computing frameworks were not designed for aerial contexts where drones must manage physical embodiment, open-world uncertainty, adversarial communication environments, and continuous fleet reconfiguration all at the same time while abiding by strict regulatory constraints. In this sense, the hardware underlying modern drones and the ambitious commercial applications built atop aerial edge intelligence have outpaced the actual capabilities of the networked computing systems needed to operate aerial fleets safely and at scale. This is the capability gap at the heart of the research agenda on aerial edge intelligence.

1.1. Motivating Scenarios

The nature of this capability gap is illustrated in the scenarios described below. These scenarios were discussed at our workshop [13] and also reflect an emerging expert consensus on the next generation of aerial applications [14].

Consider a rural hospital facing a medical emergency requiring a blood product unavailable locally. A coordinating system dispatches a heterogeneous fleet of drones optimized for speed, payload size, range, and weather conditions. The fleet uses edge AI to plan and organize the delivery mission for conditions at each stage of urgent delivery. When an unexpected storm front moves in, individual drones reroute, interrupt their plans to hand off cargo to better equipped fleet members, and renegotiate responsibilities in real time. Collectively, fleet members share situational awareness to ensure delivery within the therapeutic window. No single drone completes the mission; the fleet does, continuously adapting to conditions that static routes could not anticipate.

A second scenario involves continuous monitoring and response for early signs of wildfires. Within minutes, a network of ground sensors and citizen-contributed cameras flags anomalous heat signatures, triggering a coordinated aerial response. In this context, aerial edge intelligence dispatches small drones that fan out and characterize the fire. The data captured is streamed to firefighters, shaping the response in real time. As the fire grows, additional drones join the fleet — some carrying sensors, some acting as communication relays, some directing ground forces to staging areas, and larger security-hardened platforms delivering water or suppressant payloads. The composition of the fleet changes continuously as assets arrive, are exhausted, or are reassigned. Despite the dynamicity, the mission retains integrity, coverage is maintained, suppressant is delivered to the right locations, and the ground picture stays current. Even when the fire unexpectedly jumps a ridge and engulfs a larger area, the system does not fail; it adapts. This scenario is visually depicted in Figure 1.

These scenarios and the other applications shown in Figure 2 rely on next-generation networked systems infrastructure to enable intelligent flight planning, resource management, and data sharing as a fleet comprising heterogeneous, multimodal, and non-stationary components. In both cases, mission success depends not on the behavior of any individual drone/component but on the behavior of the fleet as a whole.

1.2. Workshop and Scope

In December 2025, we convened a workshop of 47 experts from academia, industry, and government to ask a pointed question: are the two converging trends described above sufficient in themselves to bring about aerial edge intelligence, and if not, what research, infrastructure, and policy are needed to accelerate convergence [13]? The workshop brought together researchers in networking, AI, security, robotics, and human-computer interaction alongside practitioners from the drone industry and representatives from regulatory bodies, and was supported by IEEE Computer Society and the U.S. National Science Foundation.

This paper reports the findings related to aerial edge intelligence. The framework developed reflects the authors' analysis of the workshop discussions. This is not a comprehensive summary of the workshop [13]. Our goal is to provide a shared vocabulary and a structured view of the open problems around aerial edge intelligence.

1.3. The Capability Gap

Much of the prior work on edge computing, distributed AI, federated learning, and the broader edge-intelligence agenda was created for terres-

trial, largely stationary, or single-organization settings. Aerial platforms are structurally unlike those settings. For example, a terrestrial edge node does not crash when its processor misses a scheduling deadline. A stationary sensor does not have to renegotiate its position in a communication relay while avoiding a collision. A single-organization deployment does not need to interact with an unknown peer drone mid-mission, nor terminate that interaction seconds later if the peer behaves anomalously. The first step to realizing aerial edge intelligence is to establish a framework based on such differences, shown in Figure 2, that reveals open research challenges and opportunities.

1.4. Core Properties and Challenges

Drawing on the workshop discussions, we identify five properties that together distinguish aerial edge intelligence from terrestrial edge and generic distributed computing frameworks. No existing framework fully addresses all five simultaneously:

1) **Physically embodied computation** - computation is inseparable from a mobile, energy-constrained platform that must obey the laws of physics. Unlike a terrestrial edge server, a missed deadline or a misjudged energy budget is not a performance problem; it is a safety event.
2) **Open-world, non-stationary environments** - operating conditions evolve unpredictably in space and time in ways that cannot be fully enumerated at design time, requiring systems that adapt, closing gaps in knowledge on the fly, rather than merely execute pre-specified plans.
3) **Trustworthy communication for multifaceted purposes** - wireless links must simultaneously support safety-critical control, high-bandwidth sensing data, and multi-party coordination, while resisting adversarial interference from jamming, spoofing, and man-in-the-middle attacks that terrestrial deployments rarely face at the same severity.
4) **Scalable and composable autonomous fleets** - heterogeneous agents must be added, removed, or repurposed dynamically during a mission, requiring integration principles that do not assume predictable components with fixed interfaces - precisely the assumption that AI-driven drones violate.
5) **Regulatory policy as a first-class concern** - airspace rules, privacy statutes, and safety certification requirements are not external constraints to be satisfied after the fact; they shape which problems are worth solving, which solutions are deployable, and which business models are viable. No terrestrial edge computing framework treats regulation as a first-class architectural concern.

These properties are not independent. Specific combinations give rise to three critical challenges that any aerial edge intelligence system must confront, regardless of application domain:

- **Resource-bounded runtime adaptation** - limited compute, energy, and bandwidth must be dynamically managed while the system continuously reconfigures in response to an evolving environment, with physical safety as a hard constraint on every tradeoff.
- **Ephemeral, decentralized and trusted interactions** - workable and efficacious interactions among drones, ground infrastructure, and operators must be established, maintained, and revoked without a central authority, under communication conditions that may be degraded or actively attacked.
- **Continuous compositional assurance** — safety and performance guarantees must be preserved as the fleet's membership and role assignments evolve, even though the constituent AI agents are non-deterministic and their collective behavior cannot be verified by analyzing each vehicle in isolation.

Taken together, these challenges produce a capstone problem: *emergent fleet-level behavior*. As aerial operations scale, the dominant safety concern is not the behavior of any single vehicle or any pairwise interaction, but the behavior that emerges from the fleet as a whole - behavior that is hard to infer from the design of individual agents, that was not present in simulation, and that current verification and assurance frameworks have no principled way to address.

1.5. Contributions

This paper makes the following contributions:

- We identify five core properties that together distinguish aerial edge intelligence from terrestrial edge and generic distributed computing frameworks, and argue that no existing framework simultaneously addresses all five.
- We show that these properties are interdependent: specific combinations force three critical challenges - resource-bounded runtime adaptation, trusted interactions that are ephemeral and decentralized, and continuous compositional as-

Aerial Edge Intelligence for Wildfire Detection



FIGURE 1. Aerial edge intelligence in action: Autonomous aircraft and distributed ground sensors collaborate to detect emerging wildfires, fuse uncertain observations, and maintain a shared operational picture of a rapidly changing landscape.

surance - that any aerial system must address regardless of application domain.

- We identify emergent fleet-level behavior as a capstone problem that cannot be decomposed into single-vehicle or pairwise analyses, and that becomes the dominant safety concern as aerial operations scale. We show why current verification and assurance frameworks are structurally insufficient to address it.
- Using this framework, we organize the open research challenges identified at our workshop across the areas of scaling, AI assurance, security, edge-cloud coordination, next-generation networks, and human-AI oversight, and propose a forward-looking research agenda for computing on the fly.

2. Recent Advances Toward Aerial Edge Intelligence

The convergence of two clusters of advances has made aerial edge intelligence newly tractable. The first cluster consists of advances in aerial platforms themselves: reinforcement learning for autonomous mission execution, robust localization and large-scale flight control, and improvements in battery endurance and production capacity. The second cluster consists of advances in the AI and computing infrastructure that can run on or near those platforms: resource-lean models for edge inference and planning, and data management pipelines that tie platform capability to intelligent behavior. Figure 2 places these six advances at the base of the paper's organizing framework. This section describes each in turn and explains why their convergence, while necessary, is not sufficient to deliver the commercial applications shown at the top of the figure. That gap motivates the systems framework developed in Section 3.

2.1. Advances in Aerial Platforms

Three advances have transformed drones from remotely piloted aircraft into autonomous platforms capable of executing real missions across diverse domains.

Reinforcement learning for real missions across many domains. Early drone autonomy relied on hand-coded controllers and pre-programmed routes. Reinforcement learning has changed this. Agents trained in simulation now transfer to physical platforms and execute complex maneuvers, multi-step deliveries, and coordinated search patterns across domains including agriculture, infrastructure inspection, visual object tracking, and search-and-rescue [15], [16]. This shift means that drone behavior is increasingly learned rather than specified which has significant consequences for how fleets will be verified and certified.

Robust localization and large-scale flight control. Safe autonomous flight requires that a vehicle know where it is, where its neighbors are, and how to maintain stability under disturbance. Advances in visual-inertial odometry, sensor fusion, and GPS-denied navigation have made robust localization practical on platforms that cannot carry the sensing payloads of manned aircraft [4]. Complementary work on large-scale flight control has extended these capabilities to multi-vehicle settings, enabling fleets to maintain formation, avoid collisions, and execute coordinated maneuvers without continuous human oversight [17].

Battery and flight time improvements. Power demands for drones of all sizes have decreased, and the energy density of onboard batteries has improved substantially [1], [2]. These gains directly extend mission range and payload capacity. At the same time, investment in drone hardware and production capacity has scaled dramatically, driven by both commercial demand and defense procurement [7]. The result is that capable aerial platforms are increasingly commodity infrastructure and accessible at the scale that aerial edge intelligence requires.

2.2. Advances in Edge AI

Three complementary advances have made it possible to run sophisticated intelligence on or near the con-

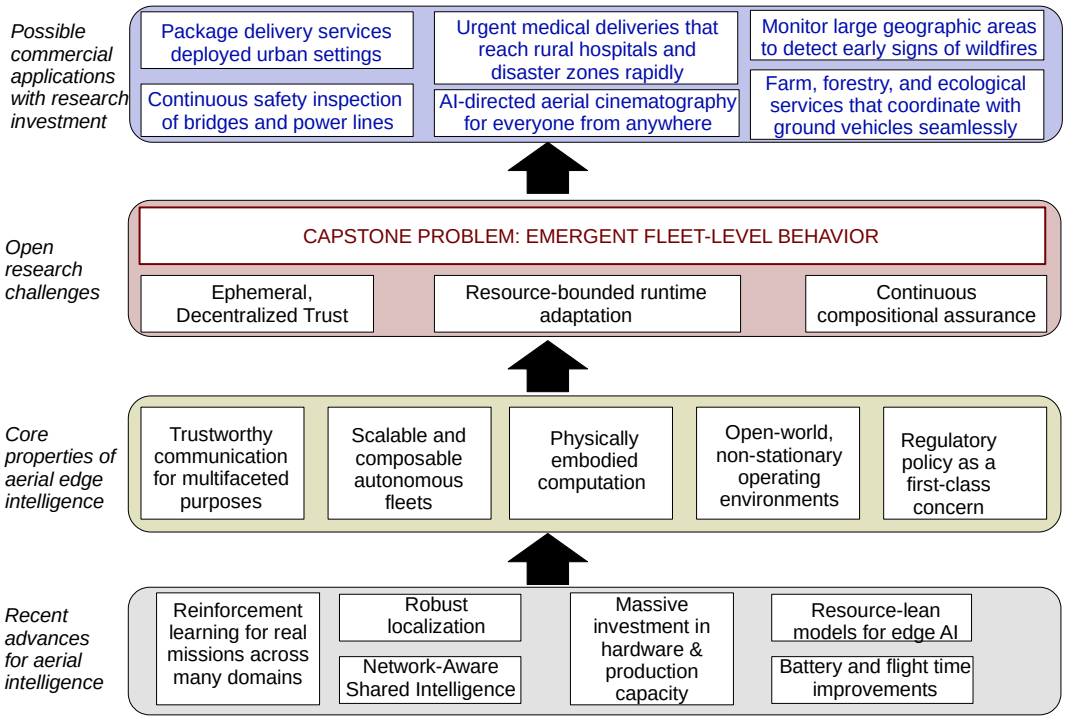


FIGURE 2. An overview of aerial intelligence from state of the art to futuristic applications.

strained platforms described above.

Resource-lean models for edge AI. The computational cost of inference and planning has fallen sharply. Model compression [8], knowledge distillation [9], and hardware-software co-design [11], [12] have produced models that achieve near-frontier performance at a fraction of the parameter count and energy cost of their cloud-hosted counterparts. These techniques are essential to host real-time AI on edge devices [18].

Network-Aware Data Management for Shared Intelligence. Edge AI advances enable drones to process sensor data locally, extract task-relevant features, and share collaborative intelligence without relying on centralized cloud processing or continuous high-bandwidth connectivity. Each drone becomes an intelligent node that filters raw sensor streams, makes local decisions, and contributes to fleet-level situational awareness in real-time.

Massive investment in hardware and production capacity. The availability of purpose-built edge AI accelerators, supported by substantial investment from both commercial and government sources, has lowered the barrier to deploying capable inference pipelines on aerial platforms [7]. This investment has also driven the development of software toolchains, simulation environments, and open datasets that reduce the cost of developing and validating edge AI systems for aerial deployment.

2.3. From Advances to Applications

Advances in aerial platforms, edge computing, and AI enable the novel commercial applications shown in Figure 2. From package delivery over vast, complex terrains to continuous safety inspection of bridges and power lines, these applications share a common requirement: intelligent, dynamic flight control that responds within milliseconds to demanding, unpredictable physical environments.

Consider wildfire detection in California's 2.3 million acres of severe fire hazard zones. Current approaches rely on satellite imagery (coarse spatial resolution, delayed detection), fixed ground sensors (immobile, expensive to deploy widely), and occasional manned flights (costly, dangerous, infrequent). Advances in aerial platforms now enable fleets of drones equipped with thermal and optical sensors to survey vast areas continuously. AI advances—particularly reinforcement learning that fuses multiple sensor modalities—allow flight controllers to make intelligent decisions about where to look, how to adapt to changing conditions, and how to coordinate with other drones. Edge computing advances bring this intelligence closer to sen-

sors, enabling drones to process data locally and share findings collaboratively across the fleet without relying on centralized cloud processing. Localization advances—particularly centimeter-level positioning and GPS-denied navigation—enable precise identification of where warning signs arise, allowing ground crews to dispatch resources to exact coordinates rather than general regions. However, wildfires spread rapidly, and early warning signs may be transient. Flight control must adjust abruptly as ground conditions change, and drones must coordinate precisely in real time.

3. Core Properties of Aerial Edge Intelligence

In this section, we highlight 5 core properties that collectively make aerial edge intelligence a unique computational workload. To be sure, the properties by themselves are not distinctive, but their combined effects produce a novel operating context.

3.1. Physically Embodied Computation

On an aerial platform, computation and flight are not separable concerns. They share the same battery, the same thermal envelope, and the same scheduling budget. A decision to run a larger model or defer a planning cycle is not a quality-of-service tradeoff; it is a tradeoff against flight time and control stability. Thermal stress and dynamic voltage scaling introduce latency jitter precisely when stability margins are thinnest. On a terrestrial edge server, a missed deadline degrades service. On an aerial platform, it can end the mission. Energy, timing, and physical safety must be treated as a single jointly constrained resource.

When computation is offloaded to nearby edge infrastructure, the constraint changes form but does not disappear. Processing delays represent unresponsive periods in a dynamic physical context: the world the drone is sensing continues to evolve while the edge server processes the previous observation. A stale perception result is not merely less accurate; in a fast-moving environment, it describes a world that no longer exists. Latency budgets for aerial edge offload are therefore dictated not by user experience expectations but by the rate at which the physical environment changes. Physically grounded world models directly address both constraints: by encoding aerodynamics, surface compliance, and mission state, they allow the drone to anticipate how the physical environment will evolve during a computation or offload cycle, reducing the cost of adaptation and making stale perception results less dangerous. Both constraints point toward the same deeper requirement: aerial edge intelligence needs world models capable of reasoning about aerodynamics, sensing context, and mission objectives as properties of the physical environment, not as abstract optimization targets [19]. A perception system that identifies a rooftop as a flat surface has done its job in a terrestrial setting. One operating on an aerial platform must also assess whether that surface can bear the drone's weight, how rotor downwash will interact with it, and whether landing there satisfies mission constraints.

3.2. Open-World, Non-Stationary Environments

The missions that motivate aerial edge intelligence are defined by their dynamism. Windy conditions degrade vision and move crops unpredictably, invalidating the plant-level health assessments that precision agriculture depends on. A herd of giraffes storming through a group of zebras dissolves the tracking assumptions a wildlife monitoring system built minutes earlier. These are not edge cases to be handled by exception logic; they are the normal operating condition. The space of situations an aerial system will encounter cannot be enumerated at design time, which means preprogrammed responses are not merely suboptimal — they can be unsafe.

This distinguishes aerial edge intelligence from terrestrial edge computing in a consequential way. A smart traffic camera can be designed around a well-characterized set of vehicle types, lighting conditions, and intersection geometries. An aerial platform operating across diverse missions — agriculture, wildlife monitoring, infrastructure inspection, emergency response — faces environments whose relevant features change not just between deployments but within a single flight. The system must reason about what it does not know, detect when its current model of the world has become unreliable, and adapt without human intervention on timescales that preclude waiting for operator guidance. Digital twins that are bidirectionally synchronized with real telemetry are a promising mechanism to simulate flight conditions without requiring exhaustive and dangerous physical flight testing, but existing frameworks do not yet capture the complexity and detail needed for aerial edge intelligence.

3.3. Trustworthy Communication for Multifaceted Purposes

Aerial platforms move at high speed through three-dimensional space, handing off between cell towers, satellite links, and peer drones. While peer drones or

higher altitude platforms can help to establish robust communication networks, it requires adaptive real-time coordination. In addition, links are unreliable and may need to operate under adversarial settings. The same channel must carry safety-critical control traffic, high-bandwidth sensor data, and multi-party coordination, with consequences of dropping any one of these that differ by orders of magnitude in severity. Beyond packet delivery, emerging airspace governance requires that every aerial platform be uniquely identifiable and trackable in flight. Identification, localization, authentication, and robust communication are coupled concerns that terrestrial networking has not needed to address together.

3.4. Scalable and Composable Autonomous Fleets

Aerial operations will scale not as a single large cluster but as many small, semi-independent fleets operating concurrently. Vehicles must join and leave each fleet at runtime due to battery exhaustion, failure, or reassignment, without degrading collective intelligence or requiring re-verification. This composability requirement is harder than single-vehicle scaling or static multi-agent coordination, both of which assume fixed membership. It is harder still because the agents are AI-driven and non-deterministic. Traditional integration assumes predictable components with well-defined interfaces. AI-driven drones violate that assumption by design.

3.5. Regulatory Policy as a First-Class Concern

Aerial failures are physical, often irreversible, and occur in shared public airspace. The technical architecture cannot be designed independently of the certification regime. AI autonomy stacks evolve continuously through retraining and operational experience, while certification frameworks assume fixed requirements and static test suites. A system whose behavior changes after certification exposes a structural mismatch with existing regulatory regimes. Treating certification as an afterthought leads either to stalled deployment or to premature operation without adequate assurance. Regulatory policy shapes which problems are worth solving, which solutions can be deployed, and which business models are viable. It must be treated as a design constraint from the outset.

4. Research Challenges

The five core properties described in Section 3 do not act independently. Specific combinations force consequences that any aerial computing system must address, regardless of application domain. We organize the field's open technical challenges under three such consequences, then show how those consequences jointly produce a capstone problem that cannot be decomposed into any one of them.

4.1. Resource-Bounded Runtime Adaptation

Open-world, non-stationary environments (Section 3.2) demand continuous runtime adaptation. But physically embodied computation (Section 3.1) means that adaptation is never free. On the drone itself, every additional computation draws from the same battery and thermal budget that keeps the vehicle airborne; a larger model or a deferred planning cycle is a tradeoff against flight stability, not just response quality. When computation is offloaded to nearby edge infrastructure, the constraint changes form: the physical world continues to evolve during the round-trip latency, so a stale result does not only reduce accuracy — it describes an environment that no longer exists. Across both settings, the world model underlying adaptation must encode not just geometric and semantic information but physical properties of the environment: aerodynamics, surface compliance, sensor context, and mission state. Adaptation mechanisms for aerial edge intelligence must therefore be cheap, anytime, physically grounded, and capable of graceful degradation.

Real-time coordination across the drone-edge-cloud continuum. Learning-based perception and planning pipelines rarely provide worst-case execution time guarantees, allowing latency and jitter to erode control stability precisely when safety margins are thinnest. Opaque execution semantics in widely used software libraries make it difficult to enforce real-time scheduling or reason about compute-energy tradeoffs in principled terms. Promising directions include autonomy frameworks that expose explicit timing, energy, and quality-of-service interfaces; standardized cross-layer reconfiguration APIs allowing drones and edge nodes to negotiate sensing rates, model fidelity, and control frequency under enforced safety constraints; and renewed attention to field-programmable hardware, which offers low-jitter, energy-efficient adaptability but remains underused due to toolchain friction with mainstream autonomy stacks.

Physically grounded perception and decision-making. Current perception pipelines suffer from what we term a depth-semantic gap: they successfully re-

cover geometry — that a surface exists — while omitting safety-critical physical semantics such as friction, compliance, and load-bearing capacity that determine whether contact with that surface is safe. Closing this gap requires sensor fusion architectures that align visual mapping with proprioceptive and inertial sensing to infer functional physical properties, and hybrid AI architectures that tightly couple learning-based perception with physics-aware world models, treating uncertainty as a first-class output rather than an afterthought.

Platform-agnostic, resource-aware autonomy. As autonomous agents, including those built on large language models, take on higher-level reasoning roles, they increasingly need to migrate across heterogeneous compute platforms — drone processors, edge servers, cloud infrastructure — as mission requirements and resource availability change. This requires execution guarantees that hold across vastly different computational and energy budgets, and resource management systems that allocate computation dynamically based on real-time constraints, mission priority, and security posture, rather than assuming a fixed deployment target.

4.2. Ephemeral, Decentralized and Trusted Interactions

Scalable and composable autonomous fleets (Section 3.4), combined with trustworthy communication for multifaceted purposes (Section 3.3), mean that interactions between drones, operators, and their constituents may require mechanisms that support identification without reliance on naive central-authority models. Fleet membership changes at runtime, and connectivity itself cannot be guaranteed, so relationships must be established dynamically and locally each time a drone joins, leaves, or reconnects — a strictly harder problem than the trust models developed for stable cloud-edge or terrestrial Internet of Things deployments.

Distributed authentication without centralized authority. Autonomous agents operating across organizational boundaries require continuous attestation during degraded or intermittent connectivity, yet current mechanisms designed for vehicle-to-everything communication remain poorly suited to this setting. Continuous dynamic attestation, in which a drone repeatedly verifies system integrity in flight without relying on a central authority, and privacy-preserving coordination protocols that allow drones to collaborate on tasks such as collision avoidance without exposing sensitive operational data, are both active and largely unsolved problems.

Next-generation geo-safety and geo-security. Unmanned aerieal vehicles (UAVs) pose significant geo-safety and geo-security risks including physical impact, airspace disruption, surveillance, and cyber-threats. Beyond-visual-line-of-sight methods include geo-fencing [20] to keep UAVs away from airports, stadiums [21], and other restricted areas. However, GPS jamming and spoofing make it difficult to honor geo-fences reliably, motivating investigation of robust positioning systems alongside scalable, geo-fence-aware three-dimensional routing and navigation methods. UAV failures and improper use further motivate exploration of safety analysis, early warning systems such as instability prediction [22], forensic infrastructure including black-box recording and failure and attack databases, and proportional countermeasures [23]. Geo-safety will also span multiple operators managing fleets in the same airspace. The operators must specify their own missions with the expectation that edge intelligence managing each drone will prevent collisions in shared, multi-authority airspace.

Identification and tracking in contested airspace. Emerging identification standards aim to make drones uniquely identifiable and trackable in flight. However, tracking methods such as radar [23] remain vulnerable to jamming and spoofing, and the autonomy stack as a whole has historically been developed with limited attention to security, creating structural weaknesses. Integrated sensing and communication, a feature anticipated in next-generation wireless standards, compounds this: the same waveform used for navigation can in principle be spoofed to fabricate obstacles or jammed to blind a vehicle, while embedding communication in sensing signals can leak operational information to passive observers. Designing identification, addressing, and tracking schemes that are robust to this dual-use vulnerability, while remaining compatible with airspace governance requirements, is an open and underexplored research area.

Securing the expanded attack surface of critical infrastructure. As ephemeral trust becomes the norm rather than the exception, the attack surface facing critical infrastructure expands accordingly: a compromised link or spoofed identity in a fleet with no persistent anchor can propagate further and is harder to attribute after the fact. Multi-modal threat detection, standardized forensic logging, and proportional, minimally collateral countermeasures are necessary complements to the authentication and identification mechanisms described above.

4.3. Continuous, Compositional Assurance

Non-stationary environments (Section 3.2) and dynamic fleet composition (Section 3.4) both demand that systems change after deployment, yet regulatory policy as a first-class concern (Section 3.5) still requires that those systems remain certified. Together, these three properties make exhaustive, one-time, pre-deployment verification infeasible. Assurance must instead become compositional — verifying subsystems and the way they combine, rather than the system as an undecomposable whole — and continuous, with assurance evidence accumulating and updating throughout operational life rather than being frozen at a single certification date.

Integrating non-deterministic agents into reliable fleets. Deep learning components may produce different outputs for identical inputs, undermining verification techniques built around reproducibility. Constraints respected during training may be violated in operational scenarios never seen during training. Coordination algorithms that behave correctly in simulation can exhibit oscillation or deadlock once deployed on real vehicles, with state spaces growing exponentially with fleet size. Addressing this requires integration testbeds spanning physical and digital environments with realistic failure injection, runtime monitoring and automated test generation for composed AI agents, and formal methods capable of scaling to multi-million-parameter models through compositional and probabilistic approaches.

Evidence-based certification for evolving systems. Static certification, built around fixed requirements and one-time test suites, is structurally mismatched with systems that legitimately evolve through retraining. Evidence-based certification frameworks, which permit iterative revision of permissible operating conditions as operational data accumulates, and which decompose autonomy into testable capabilities with assurance levels scaled to potential failure severity, offer a path forward — but require close co-design between regulators, researchers, and industry rather than purely technical solutions.

Digital twins as assurance infrastructure. Because flight testing is expensive, dangerous, and restricted to approved airspace, and because real-world failure modes cannot safely be induced during data collection, high-fidelity simulation is a necessary substrate for generating assurance evidence. This requires moving beyond simulation used solely for offline training toward digital twins with bidirectional synchronization and real-time telemetry ingestion, capable of faster-than-real-time prediction of failure modes, paired with standardized datasets of real-world failure conditions and resilience-oriented evaluation metrics that prioritize fault containment and recovery over best-case accuracy alone.

4.4. The Capstone Problem: Emergent Fleet-Level Behavior

The three challenges above do not simply sum. A fleet built from embodied, resource-bounded units (Section 4.1), held together by ephemeral and decentralized interactions (Section 4.2), and subject to continuous compositional assurance (Section 4.3), exhibits behavior at the fleet level that cannot be predicted from the analysis of any single vehicle or any pairwise interaction. Consider a fleet of delivery drones coordinating through a sudden thunderstorm: each vehicle's onboard AI makes locally reasonable decisions about path replanning, power conservation, and risk avoidance, yet the collective consequence of those individually reasonable decisions — convergence on the same narrow safe corridor, synchronized demand for charging infrastructure, cascading replanning as each vehicle reacts to the others' reactions — is a property of the fleet, not of any drone within it. Proving that such collective behavior remains safe requires verification techniques that do not yet exist at the scale the field will soon require.

This capstone problem reframes the role of scale itself. Growth from thousands to millions of aerial platforms is not simply a larger instance of the adaptation, trust, or assurance problems above; it is the condition under which emergent behavior becomes the dominant safety concern, because the combinatorial space of fleet-level interactions grows far faster than the space of individual-vehicle behaviors. It similarly reframes the role of human oversight: as a single operator comes to supervise tens or hundreds of vehicles, the task is no longer to monitor individual drones but to detect and intervene in emergent fleet-level behavior before it cascades — a supervision problem with no close analogue in single-agent or small-team human-AI interaction. Emergent fleet-level behavior is the unifying challenge that justifies treating aerial edge intelligence as a distinct systems problem, rather than as the simple union of drone-level autonomy and network-level coordination.

5. A Research Agenda for Aerial Edge Intelligence

The following research directions map to the three challenges identified in Section 4, with cross-cutting items addressing the capstone problem. Together, the core properties and their consequences organize an

foward-looking research agenda for aerial edge intelligence, structured by where contributions sit relative to the properties, the challenges they force, and the capstone problem they jointly produce.

- **Joint energy-time-safety scheduling:** Resource management frameworks that treat compute, energy, and control-stability margin as a single constrained quantity rather than as separable objectives, validated under realistic thermal and battery-contention conditions at fleet scale.
- **Compositional verification for learned autonomy:** Formal and probabilistic methods that scale to multi-million-parameter models by verifying components and their composition rather than the system as a whole, with explicit treatment of non-determinism.
- **Robust identity, addressing, tracking, and positioning for ephemeral, regulated aerial networks:** Methods for highly mobile, transient aerial nodes that are simultaneously robust and compatible with emerging airspace identification and geo-fencing requirements.
- **Digital twins as continuous assurance infrastructure:** Bidirectionally synchronized simulation environments capable of generating standardized, shareable evidence of safe behavior under open-world conditions, before and after deployment.
- **Evidence-based, capability-decomposed certification:** Regulatory and technical co-design that allows assurance levels to be granted and revised at the level of individual autonomous capabilities, rather than the system as an undecomposable whole.
- **Composable fleet architectures:** Coordination and learning architectures that admit and release individual vehicles at runtime without requiring full-fleet re-verification or degrading collective intelligence.
- **Scalable human oversight for emergent behavior:** Interfaces and supervision paradigms that surface fleet-level anomalies and uncertainty to a small number of human operators, rather than aggregating per-vehicle telemetry that does not scale.
- **Benchmarks and shared infrastructure:** Community testbeds, failure-mode datasets, and reproducible evaluation protocols spanning the physical-virtual boundary, enabling the field to compare contributions against a common baseline rather than isolated, single-institution demonstrations.

6. Conclusion

Aerial edge intelligence is more than an application of edge computing and edge AI to a new device category. Physical embodiment, open-world operation, robust mobile and identity-ambiguous networking, dynamic fleet composition, and regulatory coupling make this a challenging systems problem. This paper presented a framework that introduces the research community to the problem and details opportunities for a research agenda. We hope this framework galvanizes the research community, because moving from thousands of manually supervised aircraft toward millions of autonomous platforms operating in shared airspace will require substantial research contributions from many communities.

7. Acknowledgments

The development of this article was supported by the Computing Research Association's Computing Community Consortium through the U.S. National Science Foundation (NSF) under Grant No. 2300842. CCC Director Quinn Spadola contributed to creating the article. Any opinions, findings, conclusions, or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the U.S. National Science Foundation. Artificial intelligence tools were used to generate Figure 1 and to correct grammatical errors and writing style inconsistencies.

8. REFERENCES

1. Z. Liu, R. Sengupta, and A. Kurzhanskiy, "A power consumption model for multi-rotor small unmanned aircraft systems," in *2017 International Conference on Unmanned Aircraft Systems (ICUAS)*. IEEE, 2017.
2. B. Parveez, M. I. Kittur, I. A. Badruddin, S. Kamangar, M. Hussien, and M. Umarfarooq, "Scientific advancements in composite materials for aircraft applications: a review," *Polymers*, vol. 14, no. 22, p. 5007, 2022.
3. A. Di Sorbo, F. Zampetti, A. Visaggio, M. Di Penta, and S. Panichella, "Automated identification and qualitative characterization of safety concerns reported in uav software platforms," *ACM Transactions on Software Engineering and Methodology*, vol. 32, no. 3, 2023.
4. E. Ebeid, M. Skriver, K. H. Terkildsen, K. Jensen, and U. P. Schultz, "A survey of open-source UAV flight controllers and flight simulators," *Microprocessors and Microsystems*, vol. 61, 2018.

5. Federal Aviation Administration, “Faa aerospace forecast fiscal years 2024–2044,” U.S. Department of Transportation, Report, 2024, accessed: July 2026.
6. Federal Aviation Administration and Department of Transportation, “Normalizing unmanned aircraft systems beyond visual line of sight operations,” *Federal Register*, vol. 90, no. 2025-14992, pp. 38 212–38 391, August 2025.
7. Congressional Research Service, “DOD replicator initiative: Background and issues for congress,” Library of Congress, In Focus IF12611, January 2026, updated January 21, 2026.
8. A. O’Quinn, C. Snedeker, S. Zhang, and J. Kline, “Environment-aware dynamic pruning for pipelined edge inference,” in *2025 IEEE International Conference on Edge Computing and Communications (EDGE)*. IEEE Computer Society, 2025.
9. Y. Wang, Z. Yu, J. Wu, C. Wang, Q. Zhou, and J. Hu, “Adaptive knowledge distillation-based lightweight intelligent fault diagnosis framework in IoT edge computing,” *IEEE Internet of Things Journal*, vol. 11, no. 13, pp. 23 156–23 169, 2024.
10. S. Sridhar, S. S. Ganesh, G. Bajaj, C. Myers, and S. Parthasarathy, “Identifying knowledge gaps on the edge for visual question answering,” *Frontiers in Computer Science*, vol. 8, p. 1817034, 2026.
11. S. Heidari, M. Ghasemi, and S. Vrudhula, “Uncertainty-aware rl-based scheduling of multi-dnn workloads on edge mpsocs,” in *Proceedings of the Tenth ACM/IEEE Symposium on Edge Computing*, 2025, pp. 1–16.
12. S. Heidari, M. Ghasemi, Y. G. Kim, C.-J. Wu, and S. Vrudhula, “Camdnn: Content-aware mapping of a network of deep neural networks on edge mpsocs,” *IEEE Transactions on Computers*, vol. 71, no. 12, pp. 3191–3202, 2022.
13. C. Stewart, K. Butler, N. Aschenbruck, A. Gerall, W. Shi, D. Silver, and U. Topcu, “Computing on the fly: Navigating a vision for the future of drone computing,” Computing Community Consortium, Tech. Rep., 7 2026.
14. S. Brunswicker, J. Goppert, I. Hwang, D. Lercel, C. Song, C. Rashidian, and A. Caesar, “Airtonomy workshop report: A cyber-physical research infrastructure for next-generation autonomous aerial vehicles,” 2024.
15. J. G. Boubin, N. T. Babu, C. Stewart, J. Chumley, and S. Zhang, “Managing edge resources for fully autonomous aerial systems,” in *Proceedings of the 4th ACM/IEEE Symposium on Edge Computing*, 2019.
16. M. Bala, A. Chanana, X. Chen, Q. Dong, T. Eiszler, J. Xu, P. Pillai, and M. Satyanarayanan, “The ooda loop of cloudlet-based autonomous drones,” in *2024 IEEE/ACM Symposium on Edge Computing (SEC)*. IEEE, 2024, pp. 178–190.
17. J. Boubin, C. Burley, P. Han, B. Li, B. Porter, and C. Stewart, “MARBIE: Multi-agent reinforcement learning at the edge for digital agriculture,” in *2022 IEEE/ACM 7th symposium on edge computing (SEC)*. IEEE, 2022.
18. Q. Dong, J. Xu, P. Pillai, and M. Satyanarayanan, “Does accurate real-time AI need edge offload?” in *Proceedings of the Tenth ACM/IEEE Symposium on Edge Computing*, 2025.
19. W. Shi, Z. Dong, and P. Zhou, “Physical intelligence on the edge: A vision for the decade ahead,” *Journal of Computer Science and Technology*, pp. 1–16, 2026.
20. T. Dasu, Y. Kanza, and D. Srivastava, “Geofences in the sky: herding drones with blockchains and 5G.” New York, NY, USA: Association for Computing Machinery, 2018.
21. F. Stockman, “U.s. world cup cities are on a counter-drone spending spree,” *The New York Times*, 24th June, 2026.
22. J. J. van Beers, M. Scheffer, P. Solanki, I. A. van de Leemput, E. H. van Nes, and C. C. de Visser, “Early warning signals for loss of control in complex systems,” *Proceedings of the National Academy of Sciences*, vol. 123, no. 27, p. e2608847123, 2026.
23. G. Grieco, D. Amendola, and D. Anderson, “Detection, tracking, and identification of drones: an overview on counter-uas techniques, and open challenges,” in *2025 Integrated Communications, Navigation and Surveillance Conference (ICNS)*, 2025.