AAAI Spring Symposium Series (SSS 2026)

Toward a Closed-Loop Autonomous Sensing Framework for UAS-Based Particulate Matter Mapping

Anaiya Reliford[1], **Sonya T. Smith**[1]
[1]Howard University
Department of Mechanical Engineering
Washington, DC 20059

Abstract

Small uncrewed aerial systems (sUAS) offer a powerful capability for high-resolution characterization of fine particulate matter (PM2.5) within the atmospheric boundary layer; however, their operational utility is limited by rotor-induced aerodynamic disturbances, platform-dependent sensor bias, and the absence of integrated architectures that couple sensing, learning, and control. This paper presents an end-to-end autonomous particulate matter sensing framework that unifies sensor calibration and validation, CFD-guided sensor placement, custom payload integration, synchronized onboard computation, physics-informed measurement correction, and autonomous flight execution within a single system architecture. Computational fluid dynamics of a DJI Matrice 100 quadcopter is used to identify an aerodynamically quiescent sampling region for an Alphasense OPC-N3 optical particle counter, while a companion-computer-based pipeline synchronizes particle measurements with vehicle telemetry using MAVLink during autonomous flight. To mitigate flight-induced measurement distortion, residual-based regression models incorporating telemetry and physics-derived aerodynamic features are employed to correct airborne PM2.5 observations relative to stationary references. The resulting corrected particulate estimates provide a reliable perceptual state for spatial mapping and higher-level autonomy. Field experiments demonstrate stable autonomous data acquisition, consistent sensor–telemetry synchronization, and statistically meaningful reduction of airborne–stationary measurement discrepancies. Collectively, this framework establishes a scalable foundation for closed-loop airborne particulate matter sensing and represents a transferable autonomy pattern applicable to tactical autonomous systems that must integrate perception, learning, and control to operate under real-world uncertainty.

Introduction

Fine particulate matter (PM2.5) is a critical air pollutant associated with adverse respiratory, cardiovascular, and neurological health outcomes, motivating the need for accurate and spatially resolved measurements within the atmospheric boundary layer. Conventional air-quality monitoring networks rely predominantly on stationary ground-based instruments, which provide high-precision temporal measurements but offer limited capability to resolve three-dimensional spatial variability. As a result, localized emission sources, vertical concentration gradients, and transient pollution plumes often remain under-characterized. Small uncrewed aerial systems (sUAS) have emerged as promising platforms for addressing these limitations by enabling in situ, mobile, and altitude-resolved atmospheric sensing. Their ability to hover, execute low-speed maneuvers, and access constrained environments makes rotary-wing sUAS particularly well-suited for particulate matter sampling.

Despite this potential, several technical barriers limit the reliability and autonomy of sUAS-based particulate matter sensing. First, rotor-induced turbulence, wake interactions, and platform-generated vibrational effects distort airflow into onboard sensors, producing systematic measurement bias for optical particle counters (OPCs) mounted on multirotor platforms. Second, even when sensors are carefully integrated, airborne measurements exhibit persistent discrepancies relative to stationary reference instruments, indicating that flight-induced effects are not fully mitigated through hardware design alone. Third, many existing sUAS deployments rely on manually piloted flights or open-loop waypoint execution, decoupling environmental perception from vehicle decision-making and precluding adaptive responses to emerging concentration structures. Collectively, these limitations motivate the need for integrated architectures that jointly address aerodynamic validity, sensor calibration, mechanical integration, and synchronized onboard sensing, enabling sUAS platforms to transition from passive samplers to autonomous environmental sensing agents.

Framework Components

This paper synthesizes a series of prior studies conducted by the authors that each address a critical component of this problem. These studies include: (i) calibration and validation of low-cost optical particle counters, (ii) CFD-based identification of undisturbed sampling regions on quadcopters, (iii) design of custom 3D-printed sensor mounts, (iv) physics-informed machine-learning models for correcting airborne PM2.5 measurements, and (v) development of synchronized autonomous flight and sensing pipelines. We argue that these contributions collectively form the building blocks of a closed-loop autonomous sensing framework. This paper organizes these studies into a coherent architecture and articulates how they enable future goal-directed navigation toward pollution hotspots. This framework provides a transferable pattern for tactical autonomous systems operating under real-world uncertainty.

Sensor Calibration and Validation

Reliable autonomy begins with trustworthy perception. Prior to airborne deployment, the suitability of the Alphasense OPC-N3 optical particle counter for atmospheric research was evaluated. Three OPC-N3 sensors were subjected to field validation against a federally regulated Beta Attenuation Mass Monitor (BAM 1022) (Reliford et al. 2025). These experiments assessed PM2.5 concentration agreement as well as temperature and relative humidity performance.

The results showed good agreement for the PM2.5 measurements under static conditions, while the temperature and relative humidity exhibited systematic scaling and bias. These findings establish the OPC-N3 as suitable for atmospheric applications and motivate downstream correction strategies when operating in flight environments.

CFD-Guided Sensor Placement

To minimize aerodynamic interference with the OPC-N3 inlet, computational fluid dynamics simulations were previously performed on a simplified DJI Matrice 100 quadcopter geometry using the Reynolds-Averaged Navier–Stokes equations with an SST k–ω turbulence model. Rotors were modeled using a Multiple Reference Frame approach, and simulations were conducted across rotor speeds from 3000 to 7000 RPM under hover conditions with imposed 5 m/s headwinds and 2 m/s crosswinds (Reliford, Jackson, and Smith 2025).

Pressure, velocity magnitude, and turbulence intensity fields revealed strong downwash and wake structures beneath the rotors, as well as elevated turbulence near the arms and battery region. A relatively undisturbed flow zone was consistently identified approximately 0.2 m above the airframe. This region was selected as the target mounting location for the OPC-N3 to promote stable aspiration and reduce rotor-induced bias.

Custom Mount Design and Payload Integration

Custom 3D-printed mounts were designed to rigidly attach the OPC-N3, Raspberry Pi, and auxiliary power supply to the airframe while preserving the platform center of gravity. The printed design was made from Polylactic Acid (PLA) filament. The mount geometry maintains the OPC-N3 inlet in a vertical orientation and positions the sensor at the CFD-identified location above the airframe (Osei-Ntansah, Reliford, and Smith 2025).

Flight tests demonstrated that the modified platform maintains stable hover and maneuverability with the integrated sensor suite and that particulate and meteorological measurements collected aloft exhibit consistency with stationary observations when proper placement and mounting practices are followed.

Flight Campaign

Hover intercomparison experiments were conducted to directly compare airborne OPC-N3 measurements against a stationary reference OPC-N3. A stationary, pre-calibrated OPC-N3 was placed at an elevation of 18 m, and the DJI Matrice 100 was manually piloted to hover at the same altitude with the mounted OPC-N3 (Reliford et al. 2025).

Data were collected in Columbia, Maryland under favorable weather conditions. Both sensors recorded PM2.5 concentrations simultaneously, enabling pointwise comparison between stationary and airborne measurements. These experiments provide empirical characterization of flight-induced measurement bias and supply training data for residual-based correction models.

Physics-Informed Machine-Learning Correction

Despite optimized placement, airborne OPC-N3 measurements exhibit systematic differences relative to stationary sensors. Measurement bias is modeled as the residual error between drone-mounted and stationary PM2.5 measurements using machine learning. Both raw telemetry features and physics-derived variables were used as predictors (Walton et al. 2026).

- Telemetry features: Position, altitude, roll, pitch, yaw, linear acceleration, angular rates, three-dimensional velocity, temperature, and relative humidity.
- Physics-derived features: Total velocity, total acceleration, slip (horizontal drift) velocity, vertical oscillation speed, turbulence intensity, rotor tip speed, and estimated power from rotor-induced drag.

Multiple regression models were evaluated, including linear regression, Ridge, Lasso, random forest regression, gradient boosting, Gaussian process regression, and polynomial regression. Ridge regression produced the most stable performance, achieving coefficients of determination up to approximately 0.53 under group cross validation. This result indicates that a substantial portion of measurement bias is predictable. Importantly, this approach treats error as a learnable signal rather than random noise, producing transferable correction equations. The trained model operates as a post-processing module, transforming raw airborne PM2.5 measurements into corrected estimates.

Autonomous Flight and Sensing Synchronization

Autonomous flight control is implemented using DroneKit running on the Raspberry Pi companion computer to communicate with a CubePilot Cube Orange flight controller. DroneKit was selected due to its reliable integration with ArduPilot and consistent handling of mission commands and vehicle parameters (Osei-Ntansah, Reliford, and Smith 2026).

The OPC-N3 sampling routine executes as a Python subprocess embedded within the main autonomy script. MAVLink timestamps are used to associate each particle measurement with contemporaneous GPS position, altitude, velocity, and attitude information. This enables precise alignment of sensor outputs with vehicle state. Four classes of autonomous trajectories were employed:

- Fixed-position hover tests.
- Box (square) trajectories.
- Boustrophedon (lawnmower) coverage path.
- Horizontal zigzag profiles.

These patterns support systematic horizontal mapping, navigation precision assessment, and sensor stability analysis.

Closed-Loop Autonomous Sensing Concept

The proposed system architecture supports a closed-loop sensing-to-action pipeline that couples autonomous flight execution, synchronized onboard sensing, physics-informed correction, and navigation-level decision-making. At each control cycle, the sUAS:

- Executes an autonomous coverage or exploratory trajectory using DroneKit on the companion computer.
- Acquires OPC-N3 particulate measurements concurrently with vehicle telemetry (position, altitude, attitude, and velocity) through MAVLink-based synchronization.
- Applies a trained residual-based correction model that maps telemetry and physics-derived aerodynamic features to an estimated PM2.5 measurement bias.
- Generates corrected PM2.5 estimates and updates a spatial representation of particulate concentration.
- Modifies navigation objectives or trajectory parameters based on the inferred concentration structure.

In the present implementation, the first three steps operate online, while the remaining steps are performed offline for analysis and mapping. However, the architecture is intentionally designed to support full online integration. Because corrected PM2.5 estimates are produced onboard and time-aligned with vehicle state, they can be directly consumed by higher-level autonomy logic to influence waypoint selection, coverage density, or velocity commands.

For example, statistically significant increases in corrected concentration or sustained positive spatial gradients could trigger local trajectory refinement. This capability would enable the sUAS to autonomously increase sampling density in regions of interest and progressively localize candidate emission sources.

By embedding perception correction within the autonomy loop, rather than treating it as a post-processing step, the framework transitions the sUAS from executing static, preplanned coverage patterns to performing goal-directed environmental exploration. This establishes a foundation for scalable tactical autonomy in which aerial platforms sense, interpret, and respond to environmental structure under real-world uncertainty.

Discussion

This work reframes several previously independent studies as components of a single autonomy stack. Calibration enables trust in measurements, CFD ensures aerodynamic validity, mechanical design ensures physical realizability, autonomy software ensures repeatability, and machine learning enables perception correction. Each layer addresses a specific concern related to instrument bias, aerodynamic distortion, physical integration constraints, execution variability, and perception error, and develops an autonomous package for the detection of airborne particles.

Field experiments demonstrate reliable operation of each subsystem and validate the feasibility of the integrated architecture. The sUAS executes repeatable autonomous trajectories, OPC-N3 measurements are consistently synchronized with vehicle telemetry, and residual-based regression models produce measurable reductions in disagreement between airborne and stationary PM2.5 observations. Although correction models do not eliminate all flight-induced bias, they capture a substantial and repeatable portion of the error, indicating that a significant fraction of measurement distortion is systematic rather than random. The remaining variance suggests the presence of additional latent factors not represented in the current feature set. Ambient wind speed and direction are expected to influence slip flow, inlet aspiration, and local turbulence intensity but were not directly measured during the reported flight campaigns. Mechanical vibration of the sensor mount is another plausible contributor to variability, particularly given the limited inherent damping properties of PLA-based structures. Importantly, the observed performance demonstrates that flight-induced measurement distortion exhibits structured, learnable behavior and that physics-derived features provide meaningful explanatory power.

An important architectural implication is generality. The framework requires only three platform-level capabilities: (i) autonomous trajectory execution, (ii) access to vehicle telemetry, and (iii) companion-computer-level control of onboard sensors. As such, the architecture is not tied to a specific airframe, sensor, or autonomy library. Instead, it defines a transferable pattern for constructing autonomous environmental sensing agents.

Future Work

Future work will transition the framework from offline mapping to fully online closed-loop operation. Corrected PM2.5 estimates will be integrated into onboard decision logic to enable detection of concentration anomalies, adaptive adjustment of coverage density, and trajectory biasing toward inferred concentration gradients. This will require development of lightweight spatial representations and low-latency decision policies compatible with companion-computer execution.

To address the remaining unexplained variance in the residual correction model, future work will refine the physics-informed feature set by incorporating explicit wind-state estimation and vibration-aware metrics derived from high-rate inertial measurements colocated with the OPC-N3. These additions are intended to better capture inlet aspiration dynamics and structural effects that are not represented in the current regression framework. Model performance will be re-evaluated across expanded flight datasets to assess generalization and stability prior to full online closed-loop deployment.

Additional efforts will expand the training dataset through diverse flight conditions and environments, evaluate model generalization across platforms and sensors, and integrate real-time visualization dashboards that stream corrected measurements and vehicle state to a ground control station for supervisory awareness and optional human intervention.

Collectively, these directions aim to advance the framework from a validated autonomy architecture to an operational closed-loop environmental sensing system capable of goal-directed exploration under real world uncertainty.

Acknowledgements

The authors thank the researchers of the Applied Fluid and Thermal Engineering Laboratory (@FTERLab) in the Department of Mechanical Engineering at Howard University for their contributions to this work.

References

Osei-Ntansah, K.; Reliford, A.; and Smith, S. T. 2025. Design and Implementation of 3D Printed Sensor Mounts for sUAS-Based Atmospheric Research. In *AIAA SCITECH 2025 Forum*. Orlando, FL: American Institute of Aeronautics and Astronautics. ISBN 978-1-62410-723-8.

Osei-Ntansah, K.; Reliford, A.; and Smith, S. T. 2026. Development of Autonomous Strategies for High Resolution Data Acquisition For sUAS-Based Atmospheric Research. In *AIAA SCITECH 2026 Forum*. Orlando, FL: American Institute of Aeronautics and Astronautics. ISBN 978-1-62410-765-8.

Reliford, A.; Jackson, D.; and Smith, S. 2025. Visualizing Flow Effects on a Quadcopter to Optimize Atmospheric Sensor Placement. In *AIAA AVIATION FORUM AND ASCEND 2025*. Las Vegas, Nevada: American Institute of Aeronautics and Astronautics. ISBN 978-1-62410-738-2.

Reliford, A.; Smith, S. T.; Osei-Ntansah, K.; Stewart, A.; and Jackson, D. 2025. From Calibration to Flight: Assessing the Alphasense OPC-N3 for Aerial Data Collection. In *AIAA SCITECH 2025 Forum*. Orlando, FL: American Institute of Aeronautics and Astronautics. ISBN 978-1-62410-723-8.

Walton, N.; Reliford, A.; Cooper, C.; and Smith, S. T. 2026. Examining Machine-Learning Methods to Develop a Correction Equation for Drone-Mounted PM2.5 Concentrations. In *AIAA SCITECH 2026 Forum*. Orlando, FL: American Institute of Aeronautics and Astronautics. ISBN 978-1-62410-765-8.