

THREE-DIMENSIONAL ADAPTIVE TRAJECTORY OPTIMIZATION FOR RPAS-BASED ATMOSPHERIC MEASUREMENTS

Kuno A. Buchtal[1] & Alexander Lau[1]

[1]German Aerospace Center (DLR), Institute of Air Transport, Hamburg, Germany

Abstract

Airborne contamination such as volcanic ash and radioactivity threatens aviation safety, yet in-situ measurements that could constrain the forecasts are sparse, and an aircraft collecting them must decide where to fly while the situation evolves. We address this with a real-time, adaptive three-dimensional trajectory optimization for a remotely piloted aircraft system that plans its route to be maximally informative about the contamination field and re-plans in flight as measurements arrive. Extending earlier two-dimensional, static work, the system couples an operational data chain into a closed loop: a numerical weather prediction from the German Weather Service initializes a three-dimensional field, a genetic algorithm optimizer at the German Aerospace Center plans a flight-ready route, and an in-situ measurement payload operated by enviscope returns observations that are blended back into a working field to drive in-flight re-optimization. A lightweight linear interpolation replaces the computationally expensive estimator of the earlier work, so each candidate route can be evaluated within the real-time loop. The system was deployed in a 2025 flight campaign in Andøya, Norway, on a synthetic surrogate field. On a controlled field the optimizer converges reliably, and in flight the full chain executed end to end: optimized trajectories were flown, incoming measurements were merged into the working field, and a re-optimized trajectory was computed, uploaded and accepted by the flight operator. We close by noting where a purely exploitative objective becomes uninformative, which points to an exploration term as the next step.

Keywords: 3D trajectory optimization, RPAS, atmospheric measurements, flight validation, real-time adaptation, genetic algorithm

1. Introduction

The 2010 eruption of Eyjafjallajökull in Iceland stranded more than 10 million passengers and caused airline revenue losses exceeding $1.7 billion [1], exposing a substantial gap in airborne hazard monitoring for aviation safety [2].

The collaborative project MEASURE, funded by the German Federal Aeronautical Research Programme (LuFo VI-1, FKZ 20V1906B) and led by the German Weather Service (DWD), develops, integrates and verifies a remotely piloted aircraft system for in-situ airborne measurement of particulate and radiological data in contaminated airspace. The work divides along clear roles. DWD leads the project and provides the meteorological and contamination forecasts. DLR develops the three-dimensional trajectory optimization that plans the flight path from the incoming measurements. enviscope GmbH builds, integrates and operates the measurement payload, which runs on its own board computer and data link and is technically decoupled from the platform. The platform itself is the Airbus Defence and Space DO-TP25 (Figure 1, Table 1), selected to meet strict mission requirements for payload, range and operating altitude.

For the MEASURE campaigns the DO-TP25 carries a board computer and instruments in a nose compartment, together with interchangeable wing-pod payloads for the two target hazards: (1) optical

https://doi.org/10.71945/icas2026_0292

particle counters and filter collectors for volcanic ash, and (2) gamma spectrometry, ambient-dose-rate probes and filter sampling for radioactivity. Differently equipped deployable dropsondes are used for vertical profiling.
The trajectory optimization is independent of the specific payload, since every sensor enters the chain as a scalar measurement on the grid. As no volcanic-ash or radioactivity event occurred during the campaign, and no such hazard could be measured in a test case, the present study uses relative humidity as a surrogate. It is a quantity that the payload measures directly in clear air, and it exercises the full optimization and merging loop in exactly the same way a contamination measurement would. In the final envisioned operational use case, DWD's ICON-ART (ICOsahedral Nonhydrostatic – Aerosols and Reactive Trace gases) numerical weather prediction (NWP) model [3] forecasts contamination propagation and provides the initial concentration fields.



Figure 1 – DO-TP25 remotely piloted aircraft system (RPAS) with measurement equipment in nose compartment and wing pods during pre-flight preparation.

At ICAS 2024 [4], we presented a genetic algorithm (GA) based two-dimensional route optimization using Kriging estimation on ICON-ART forecast data. Kriging reconstructed an estimated contamination field from the values sampled along a candidate route, and the algorithm optimized the waypoint set to minimize the mean absolute error (MAE) between that estimate and the reference field. However, that approach suffered from three principal limitations: it neglected the vertical dimension, it supported only static pre-flight planning without the capability to adapt to live data, and its reliance on Kriging proved to be too computationally expensive for fast, in-flight three-dimensional re-optimization loops.
The present paper addresses these limitations. Our core contributions are: (i) a fully three-dimensional spatial representation and trajectory optimization framework; (ii) a computationally lightweight estimation stage based on linear interpolation that replaces Kriging to enable real-time evaluation; (iii) a closed-loop, in-flight re-optimization architecture driven by a spatially weighted measurement-merging product; and (iv) the operational demonstration of the complete data chain, from numerical weather prediction to an approved in-flight trajectory, during a 2025 flight campaign using synthetic data that replicated the anisotropic nature and granularity of operational weather fields.

2. Related Work

Informative path planning (IPP) seeks trajectories that maximize the information gathered about an environment under motion and budget constraints, and it is a natural framework for environmental monitoring with sensor-equipped aircraft [5]. Because such missions must balance competing aims, much of the literature treats IPP as multi-objective, weighing travel time and energy against coverage and information gain [6, 7]. In disaster response in particular, adaptive in-situ sampling has been

Table 1 – DO-TP25 flight performance data.

Performance parameter	Value @ SL, ISA
Maximum speed	300 kts (154 m/s)
Endurance	60 min (mixed speed)
Operating altitude	35 – 30,000 ft (10 – 9144 m)
Operating range	110 km Line of Sight
Payload capacity	25 kg

shown to improve both the speed and the accuracy of situational assessment [8].
GAs are well suited to this setting, because aerial path planning is non-convex, multi-objective and heavily constrained, and GAs remain robust across offline and online as well as two- and three-dimensional formulations [9, 10]. The two regimes serve different purposes: Offline planning computes a route before the mission for well-mapped, predictable environments [11], whereas online planning adapts to changes as they are observed during flight [12]. The present work applies a genetic algorithm in both regimes. The same optimizer and parameters are used pre-flight and in flight; during the mission it re-plans only the remaining route, from the current aircraft position and against the updated field.
At the core of adaptive sampling lies the trade-off between exploitation and exploration. Exploitation refines the estimate where structure is already visible, while exploration samples where the current estimate is most uncertain. The canonical way to balance the two is the Gaussian-process upper confidence bound (GP-UCB). It scores each candidate point by combining the current estimate with the model's uncertainty at that point, $\mu(x)+\beta\,\sigma(x)$, where the posterior mean $\mu(x)$ drives exploitation, the posterior uncertainty $\sigma(x)$ drives exploration, and β sets the balance between them [13]. The decisive ingredient is the uncertainty $\sigma(x)$: without it there is no exploration term to add, so the whole approach depends on having an estimator that provides one. The same principle drives online IPP for environmental monitoring, where a Gaussian-process belief steers sampling toward uncertain regions [14, 15]. Its drawback is cost. Computing the Gaussian-process variance scales at least quadratically with the number of observations, which becomes prohibitive when an optimizer must evaluate many candidate paths per iteration [13]. This tension, between a principled uncertainty signal and real-time feasibility, is exactly the gap addressed in Section 4.6.
Most IPP and GA studies are validated in simulation alone. The distinguishing feature of the present work is that the complete chain was exercised in real flight, from forecast data through three-dimensional optimization, flight-ready export and flight operator acceptance, and including an in-flight re-optimization computed and uploaded during a live mission.

3. System Architecture

3.1 Operational Data Chain

The system spans three domains, shown in Figure 2: a forecast domain (DWD), an optimization domain (DLR), and a field domain that holds the ground stations and the RPAS operated by enviscope GmbH and Airbus D&S.
The forecast domain provides the contamination scenario. In the operational target configuration this is the ICON-ART composition model, which propagates aerosol and trace-substance fields [3]. As no contamination event occurred during the 2025 campaign, a synthetic field was used in its place (Section 3.2). A interface for real pressure-level NWP data has also been implemented for the extended campaign (Section 7.).

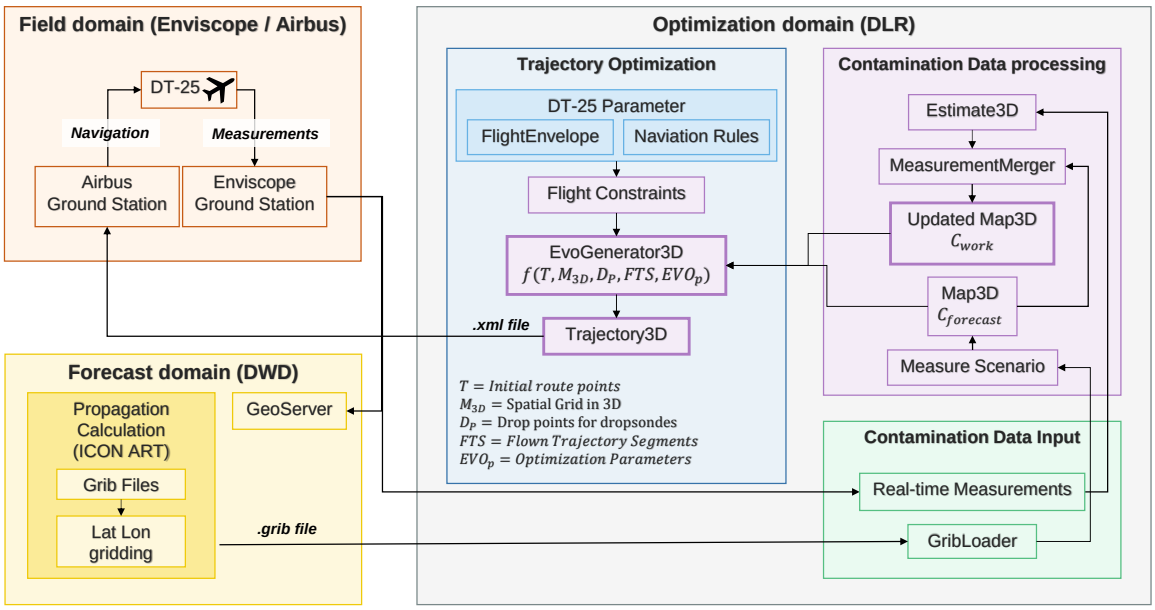


Figure 2 – System architecture showing DWD forecast generation and DLR flight planning components with real-time measurement feedback loop.

The DLR optimization domain holds three components. The first is the Contamination Data Input. Its `GribLoader` ingests the forecast, or the surrogate, and reads the MEASURE scenario. A contamination-data processing step then translates this into the internal `Map3D` grid, which carries the scalar field C and is initialised to the forecast C_{forecast}. The same component also handles the live data. Incoming measurements are turned into a field estimate by `Estimate3D`, and the `MeasurementMerger` blends that estimate into the grid, updating it to the working field C_{work} (Section 4.5). The second component is the Trajectory Optimization. It reads the current `Map3D` together with the flight constraints, which combine the `FlightEnvelope` and the `NavigationRules` with the DO-TP25 platform parameters, and runs the `EvoGenerator3D` to produce a `Trajectory3D`. The third step exports that trajectory and hands it to the field domain.

The field domain comprises two separate ground stations and the RPAS, operating on distinct data links. The Airbus D&S ground station carries the navigation and control network. It uploads the trajectory data to the RPAS and is used by the operator and mission control to fly the platform, either manually or in an automated manner. The ground station is platform-agnostic and decoupled from the control link. It receives the atmospheric measurements, from both the RPAS and the dropsondes, over the LoRa data link, and forwards them by SFTP to DWD and DLR servers. On DWD side these data support real-time analysis and visualization for flight operators and aviation decision-makers. On the DLR side they feed the real-time measurement input described above. This paper concerns DLR side, that is, the update of the working field and the re-optimization it drives. In the 2025 flights the in-situ sensors sampled at 1 s, and the data link delivered the measurements in one-minute batches as separate CSV files to the incoming servers.

The architecture is built for more than the RPAS track. It accepts up to eight dropsondes, whose additional vertical profiles enter the same estimation and merging stages as point observations. This paper reports results for the RPAS in-situ track alone.

3.2 Synthetic Scalar Field

For a controlled study of the optimizer in isolation from data-pipeline effects, a synthetic three-dimensional scalar field, the *sample map*, is used as the reference field. It is a single three-dimensional Gaussian plume, intended to resemble the data *structure* of an ICON-ART contamination forecast for a volcanic-ash scenario such as in [4], not to reproduce its physical content. In the convergence

study of Section 5. this field represents the reference C_{ref} that the optimizer is scored against, with pseudo-measurements drawn from it along each candidate trajectory. In the Flight 2 merge it is the forecast prior C_{forecast} into which the real measurements are blended. A separate offset truth field, which a twin experiment would require in order to evaluate MAE_{val} against MAE_{base}, is not used here; that evaluation is reserved for the extended campaign (Section 7.).

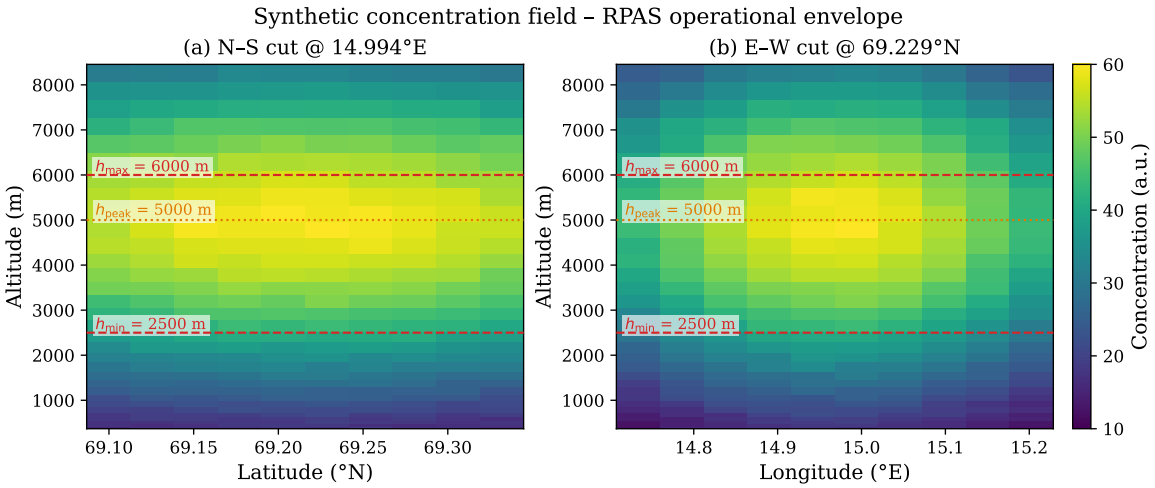


Figure 3 – Synthetic sample map (Andøya configuration): (a) east–west vertical slice at peak latitude $69.216°$N, (b) north–south vertical slice at peak longitude $14.968°$E, (c) horizontal slice at the model level closest to the peak altitude of 5000 m. The dashed line marks the peak altitude; the non-equidistant vertical levels and the horizontal grid anisotropy are visible. $h_{\min}$ and $h_{\max}$ are the minimum and maximum altitude for the waypoint generation.

It is intended to resemble the data *structure* of an ICON-ART contamination forecast for a volcanic-ash scenario such as in [4], not to reproduce its physical content. A peak of value 1000 (arbitrary units) is placed at the geometric centre of the Andøya bounding box ($14.968°$E, $69.216°$N, 5000 m). The plume width is set by horizontal standard deviations $\sigma_{\text{lon}} = \sigma_{\text{lat}} = 0.3°$ and a vertical standard deviation $\sigma_z = 3000$ m. Additive Gaussian noise with a standard deviation of 10 is then applied, and the field is rescaled to $[10, 60]$. The horizontal grid is 10×10 over the box $[14.734°, 15.203°]$E $\times$ $[69.100°, 69.332°]$N, giving cell spacings of approximately 2058 m east–west and 2871 m north–south at the Andøya latitude. ICON-EU, the operationally relevant product at this latitude, has a native resolution of $0.0625°$, roughly 2469×6958 m. The two are of the same order, and close in the east–west spacing. Vertically, the grid uses 33 levels from 378 to 8257 m, with non-equidistant spacing from approximately 20 m near the surface to 395 m aloft. Figure 3 shows two vertical cross-sections through the peak and a horizontal slice at peak altitude.

4. Methodology

4.1 Spatial Representation and Operational Corridor

The atmosphere is represented by a three-dimensional grid (`Map3D`) as described above. Its vertical levels are derived from the ICON model levels, which are terrain-following; here they are simplified to geometric heights above mean sea level by horizontal averaging, so the levels are flat across the grid. Each grid point holds one scalar value of the field C, a contamination concentration in the target application or a meteorological surrogate in the present study. All estimation and error computations are performed on this grid. The trajectory search is restricted to an *operational corridor* defined by the configurable minimum and maximum waypoint altitudes ($h_{\min}$ and $h_{\max}$) (as shown in Figure 3), which in the Andøya configuration were set to 2500 m and 6000 m, respectively. This corridor reflects the mission profile of the RPAS, and it is distinct from the underlying `Map3D` grid, which extends well beyond it.

4.2 Spatial Estimation

Given the set of measurements taken along a candidate trajectory, a field estimate $\hat{C}$ is reconstructed over the full grid. The estimator (`FastEstimate3D`) supports several methods, of which linear interpolation and radial-basis-function (RBF) interpolation represent practicable options. Linear interpolation was used throughout the campaign. It has no shape parameter to tune, it is robust on scattered samples, it is monotone between data points and so does not overshoot, and it is the cheapest of the available methods. This is a deliberate departure from the Kriging estimator of [4]. Kriging provides a posterior variance natively, but it is too expensive per evaluation for in-flight use, where the estimator is called once per individual per generation. Linear interpolation is computationally efficient, but it provides no estimate of estimation variance. This limitation directly motivates the exploration strategy introduced in Section 4.6.

4.3 Genetic Algorithm and Fitness

A trajectory is encoded as an ordered set of three-dimensional waypoints, each carrying a latitude, longitude and height plus a navigation transition type. The genetic algorithm itself follows our two-dimensional formulation [4], with the same population structure, selection, crossover, mutation and elitism. Three changes adapt it to the present problem. First, the waypoints and the search space are three-dimensional and confined to the operational corridor. Second, every candidate route is checked against the flight envelope and is either repaired or rejected (Section 4.4). Third, the spatial estimate uses linear interpolation in place of Kriging, which removes the per-evaluation cost that made the previous estimator too slow for in-flight use. Table 2 lists the parameters. The same parameters are used pre-flight and in flight; only the planning horizon shrinks during the mission, since part of the route is already flown.

The fitness rewards a trajectory for being informative about the field. The key idea is that the optimizer does not need real measurements to score a route. For each candidate trajectory it reads the field values at the points the aircraft would sample along that route, and reconstructs a full field estimate $\hat{C}$ from those samples by interpolation (Section 4.2). A good trajectory is one whose samples are sufficient to reconstruct the field well; a poor one leaves large regions unconstrained. The fitness measures this directly, as the mean absolute error between the reconstruction $\hat{C}$ and the reference field C_{ref} over all grid points,

$$\mathrm{MAE}_{\mathrm{fit}} \;=\; \mathrm{MAE}\big(\hat{C}, C_{\mathrm{ref}}\big) \;=\; \frac{1}{N}\sum_{i=1}^{N}\big|\hat{C}_i - C_{\mathrm{ref},i}\big|, \tag{1}$$

where N is the number of grid points and the GA minimizes $\mathrm{MAE}_{\mathrm{fit}}$.

Three fields must be kept apart here. The estimate $\hat{C}$ is what the route reconstructs from its own samples. The reference field C_{ref} is whatever field the route is currently scored against: before the flight this is the forecast C_{forecast}, the externally provided prediction that initializes the grid (the ICON-ART composition forecast, or the synthetic surrogate in the 2025 campaign), and it is not derived from any measurement; during the mission it is the working field C_{work}, the forecast after real measurements have been merged into it (Section 4.5). The fitness is the estimation error alone. The three path strategies, shortest, descending and energy-efficient, only set the order in which a given waypoint set is visited; they produce a tour, not a fitness contribution.

4.4 Constraints and Navigation

Candidate trajectories must respect flight-dynamic feasibility. The FlightEnvelope module provides the DO-TP25 performance limits. Climb rate, descent rate, airspeed bounds and turn radius are interpolated from lookup tables as functions of altitude and airspeed. Each trajectory segment is validated against

$$h_{\min} \le h(t) \le h_{\max}, \qquad \dot{h} \le V\,\tan\big(\gamma_{\max}(h,V)\big), \qquad V_{\min}(h) \le V \le V_{\max}(h), \tag{2}$$

where h is altitude, $\dot{h}$ the climb or descent rate, V the airspeed and $\gamma_{\max}$ the maximum climb or descent angle from the envelope. The turn radius is taken from the envelope as an altitude-dependent value

Table 2 – Parameters of the three-dimensional genetic algorithm (`EvoGenerator3D`).

Symbol	Description
T	Initial route points (support waypoints)
C_{ref}	Reference field on the `Map3D` grid
DP	Drop points for dropsondes (optional, default 0)
FTS	Flown trajectory segments (fixed prefix in re-optimization)
N_{pop}	Population size
G	Number of generations
P_c	Crossover rate
P_m	Mutation rate
E	Number of elites carried to the next generation
Selection	Tournament or rank, with elitism
Strategy	Path ordering: shortest, descending or energy-efficient
Handling	Constraint handling: repair or reject

$R_{\text{turn}}(h)$. Constraints are enforced by a repair-or-reject mechanism rather than by an additive penalty in the fitness. An optional penalty mode exists but is not the default.
Navigation-aware generation implements standard waypoint handling. It supports *FlyBy* and *FlyOver* waypoint types and *TrackToFix* and *DirectToFix* leg construction. *FlyOver* passes through the waypoint before turning. *FlyBy* starts the turn ahead of the waypoint.

4.5 Real-Time Adaptive Re-optimization

During flight, incoming measurements are merged into the forecast to form the working field, by a spatially weighted blending scheme (`MeasurementMerger`). The forecast here is the same externally provided field C_{forecast} that initializes the grid (Section 4.3); the merge updates it into the working field C_{work} as real measurements arrive. All available samples are used together, the RPAS in-situ track and any dropsonde profiles. They are first interpolated onto the forecast grid by a single linear interpolation, which yields a measurement-derived field $\hat{C}_{\text{meas}}$. A spatial weight $\alpha(x)$ then controls, cell by cell, how far the analysis moves from the forecast C_{forecast} toward $\hat{C}_{\text{meas}}$.
For each grid point x the weight is based on the *nearest* measurement, not on a sum over all measurements,

$$\alpha(x) \;=\; \exp\!\left(-\frac{d_h^{\min}(x)^2}{2\sigma_h^2}\right)\cdot\exp\!\left(-\frac{d_v^{\min}(x)^2}{2\sigma_v^2}\right), \tag{3}$$

where $d_h^{\min}$ is the geodesic horizontal distance and $d_v^{\min}$ the vertical distance to the closest measurement. The falloff is Gaussian, with separate horizontal and vertical length scales σ_h and σ_v. The two scales differ because atmospheric fields vary faster vertically than horizontally. A measurement informs a wide horizontal neighborhood but only a thin vertical layer. In the campaign configuration $\sigma_h = 10\,\text{km}$ and $\sigma_v = 500\,\text{m}$. These scales are of the order of atmospheric correlation length scales, which are tens of kilometres horizontally for moisture and much shorter vertically, of order hundreds of metres in the boundary layer [16]. The chosen values sit at the local end of those ranges, which spreads sparse in-situ samples without over-smoothing. A formal tuning against a known truth is future work; the sensitivity to the two scales is shown in Section 5. (Figure 11). The weight is then clipped to $[0, 1-w_f]$, so the forecast keeps a minimum share w_f even at a measurement location. With $w_f = 0.1$ the analysis retains at least ten percent forecast everywhere, which keeps the forecast as a stabilizing prior and prevents a single outlier from overwriting the field. The working field is

$$C_{\text{work}} \;=\; \alpha\,\hat{C}_{\text{meas}} \;+\; (1-\alpha)\,C_{\text{forecast}}. \tag{4}$$

This background-plus-observations blending is a pragmatic, single-pass form of successive-corrections objective analysis, the family introduced by Cressman [17] and Barnes [18], in which a first-guess field is nudged toward observations by distance-dependent weights. It is deliberately simpler than optimal interpolation or variational assimilation, which would additionally require background and observation error covariances, and it is chosen here for its low cost inside the real-time loop.

Because the weight uses only the nearest measurement and the interpolation is local, the scheme is far cheaper than a global covariance model such as Kriging (Section 4.6). The blending is purely spatial. Temporal weighting is reserved as a future extension.

Re-optimization is triggered on a time basis rather than on a field-divergence test. Two conditions must hold. First, the flight time must exceed an initial delay t_{delay}, which suppresses re-optimization during the climb, when only a thin and changing slice of the atmosphere has been sampled. Second, the time since the last re-optimization must exceed a minimum interval t_{min}, which prevents re-planning on every cycle. In the campaign configuration $t_{\text{delay}} \approx 600\,\text{s}$ and $t_{\text{min}} \approx 180\,\text{s}$. When both hold, the loop merges the new measurements, re-runs the optimizer against C_{work}, and uploads the resulting trajectory. A manual trigger was added after the 2025 campaign and is not used in the results reported here. Algorithm 1 summarizes the loop. At re-optimization the already flown trajectory segments are passed to the optimizer as a fixed prefix, so the new plan continues from the current aircraft position rather than re-planning the whole route. The same estimation and merging stages accept dropsonde profiles as additional point observations, although the results below use the RPAS in-situ track alone.

Algorithm 1 Real-time adaptive re-optimization

$C_{\text{work}} \leftarrow C_{\text{forecast}}$ {`Map3D` initialized with the forecast field}
plan $\leftarrow$ EVOGENERATOR3D(C_{forecast}) {pre-flight optimization}
$t_{\text{last}} \leftarrow 0$
while mission active **do**
  $M \leftarrow$ ACQUIRENEWMEASUREMENTS() {read newly arrived CSV batches from the data link}
  if $M \neq \emptyset$ **then**
    $C_{\text{work}} \leftarrow$ MERGE(C_{forecast}, M) {Eq. 3,4}
  end if
  if $t_{\text{flight}} \geq t_{\text{delay}}$ **and** $t_{\text{flight}} - t_{\text{last}} \geq t_{\text{min}}$ **then**
    plan $\leftarrow$ EVOGENERATOR3D(C_{work}) {re-optimize, minimize Eq. 1}
    EXPORTANDUPLOAD(plan)
    $t_{\text{last}} \leftarrow t_{\text{flight}}$
  end if
end while

The loop runs once per polling cycle. ACQUIRENEWMEASUREMENTS checks the incoming directory for measurement files that have arrived over the data link and reads them. MERGE blends those measurements into the forecast to update C_{work}, as in Eq. 3 and Eq. 4. When the two timing conditions hold, EVOGENERATOR3D re-plans against C_{work} and the new trajectory is exported and uploaded.

4.6 Exploration Term (Proposed Extension)

The fitness of Eq. 1 is purely exploitative. It rewards a trajectory whose reconstruction matches the reference field. When the field is spatially homogeneous, MAE_{fit} is almost constant across candidate trajectories, so the objective gives essentially no optimization direction. A reduction of MAE_{fit} presupposes exploitable spatial variation. This is a structural property of the objective, independent of any particular field, and it matters wherever the true field is locally smooth.

As a principled extension, presented here as future work and not as an implemented result, we propose adding an exploration term,

$$\text{Fitness} = \text{MAE}_{\text{fit}} - \lambda\, U(\text{traj}), \tag{5}$$

where U rewards sampling in under-observed regions and λ balances exploration against exploitation. The natural choice for U is a posterior uncertainty, since a trajectory that visits high-variance cells is expected to be informative. That signal is exactly what our estimator does not provide. Linear interpolation returns no variance. Kriging and Gaussian-process regression return one natively, but they are the same method and share the same cost. Building the covariance matrix over n measurements requires $O(n^2)$ memory and solving it requires $O(n^3)$ time, where n grows with the number of

in-situ samples and dropsonde points. In our trials Kriging exhausted available memory at the sample counts produced by a single flight, before dropsondes were even added, which is why it is infeasible as a per-evaluation step inside the genetic algorithm. We therefore propose a cheap surrogate that needs no covariance matrix. A visit-count bonus maintains a counter per grid cell and rewards a candidate trajectory for entering cells that have been sampled little or not at all, in the spirit of count-based exploration. This is constant cost per cell. The dropsonde profiles fit naturally here, since they provide the vertical coverage that the largely horizontal RPAS track under-samples.

5. Results

We distinguish three error quantities, all defined on the `Map3D` grid and all built on the three fields introduced in Section 4.3 ($\hat{C}$, C_{forecast} and C_{work}, with C_{truth} added in simulation). They answer different questions, and not all are observable in flight.

- $\text{MAE}_{\text{fit}} = \text{MAE}(\hat{C}, C_{\text{ref}})$ is the optimizer-internal self-consistency, the quantity the GA minimizes (Eq. 1). It needs no ground truth, so it is the only error observable in real flight. It measures how well the trajectory reconstructs the field it is optimized against, not accuracy against the real atmosphere.
- $\text{MAE}_{\text{val}} = \text{MAE}(C_{\text{work}}, C_{\text{truth}})$ is the accuracy of the analysis against a known truth field.
- $\text{MAE}_{\text{base}} = \text{MAE}(C_{\text{forecast}}, C_{\text{truth}})$ is the forecast-only baseline.

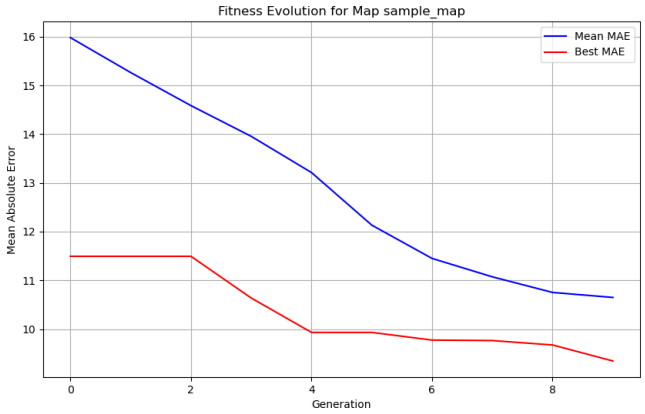


Figure 4 – Convergence of the optimizer objective MAE_{fit}, best and mean, over generations on the synthetic field, campaign configuration (10 generations).

The statement that the system improves the estimate is $\text{MAE}_{\text{val}} < \text{MAE}_{\text{base}}$, that is, the measurement-informed analysis is closer to the truth than the forecast alone. Both quantities require a known truth field and can therefore be evaluated only in a twin experiment. Such a controlled comparison is not part of this paper. It is established for the extended campaign, which uses real ICON-D2 fields as truth (Section 7.). This paper is a feasibility demonstration. It reports the operational execution of the chain and the convergence of MAE_{fit}, and it does not present the flight convergence as accuracy against the real atmosphere.

5.1 Algorithmic Convergence on the Synthetic Field

On the synthetic sample map of Section 3.2, the GA was run with the campaign configuration. Figure 4 shows MAE_{fit}, best and population mean, over generations. The objective decreases steadily and the best individual approaches a plateau. With these settings convergence is typically reached near generation 15. The campaign run was limited to 10 generations, because the on-site compute ran on a single laptop and the connection to the data server was restricted, so the curve is shown to that point.

5.2 Flight Campaign 2025 (Andøya)

Two validation flights were conducted during the 2025 flight campaign in Andøya, Norway. Both used the same configuration: a population of 100 over 10 generations, 12 initial waypoints, the operational corridor 2500 to $6000\,\mathrm{m}$, and linear interpolation. The surrogate field was the synthetic sample map of Section 3.2, scaled to the range $[10, 60]$ so that relative-humidity measurements could be merged directly.

Flight 1: trajectory execution Flight 1 covered the 12 planned waypoints across the operational area at altitudes between approximately 4500 and $6000\,\mathrm{m}$. The end-to-end function was demonstrated, from trajectory generation through waypoint navigation with continuous measurement acquisition. The navigation procedures, TrackToFix and FlyOver, were correctly interpreted and executed by the RPAS after the operator check at the ground station. Two in-flight decisions by the flight op-

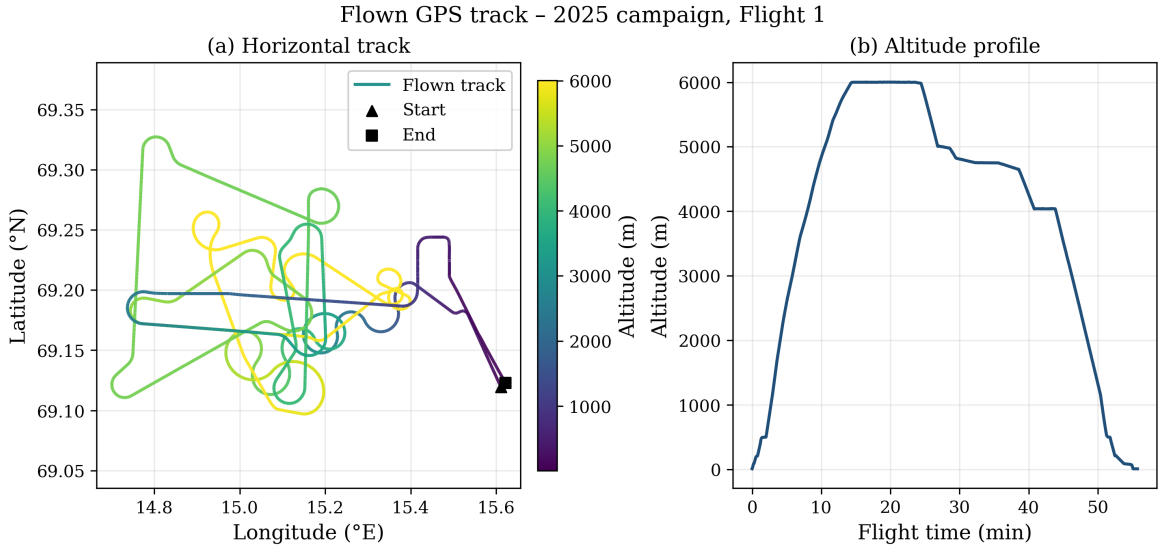


Figure 5 – Reconstructed GPS track of the 2025 Andøya campaign, Flight 1: horizontal track colored by altitude (left) and altitude-over-time profile (right). Position was reconstructed from the Airbus flight-controller GPS, as the enviscope computer recorded none.

erator changed the plan. Waypoint 8 was skipped because a ship lay beneath it, and waypoint 6 was moved slightly for the same reason. Waypoints 11 and 12 were then skipped near the end, to preserve flight time and ensure a safe return, after the waypoint-following objective had been met. The data-acquisition computer recorded no GPS on this flight, so the measurement positions were reconstructed afterwards by interpolation onto the Airbus flight-controller GPS track, which did record position. For the same reason no in-flight merge or re-optimization was attempted on Flight 1, since real-time measurement positions were not available. Figure 5 shows the reconstructed track, colored by altitude, with the altitude-over-time profile, and Figure 6 compares the planned and flown waypoints.

Flight 2: in-flight merge and re-optimization
Flight 2 targeted the closed loop and did record measurement-side position, which enabled the merge. After an initial phase of hardware and integration checks, the aircraft flew the optimized trajectory (Figure 7). Between minutes 32 and 34 after launch it reached two waypoints of that trajectory, as planned (Figure 8). Due to technical reasons on the platform side no further waypoints were reached after the second waypoint. The incoming relative-humidity measurements were merged into the forecast field with the scheme of Section 4.5, producing the working field C_{work}, and a re-optimization was computed in flight, exported and uploaded, and accepted by the flight operator.
For this single configuration the operational timing was measured directly. This is not a parameter study, only a characterization of the setup that was flown. The optimization, with a population of 100 over 10 generations and 12 initial waypoints (compare Table 2), took $98\,\mathrm{s}$ with 10 processes in parallel.

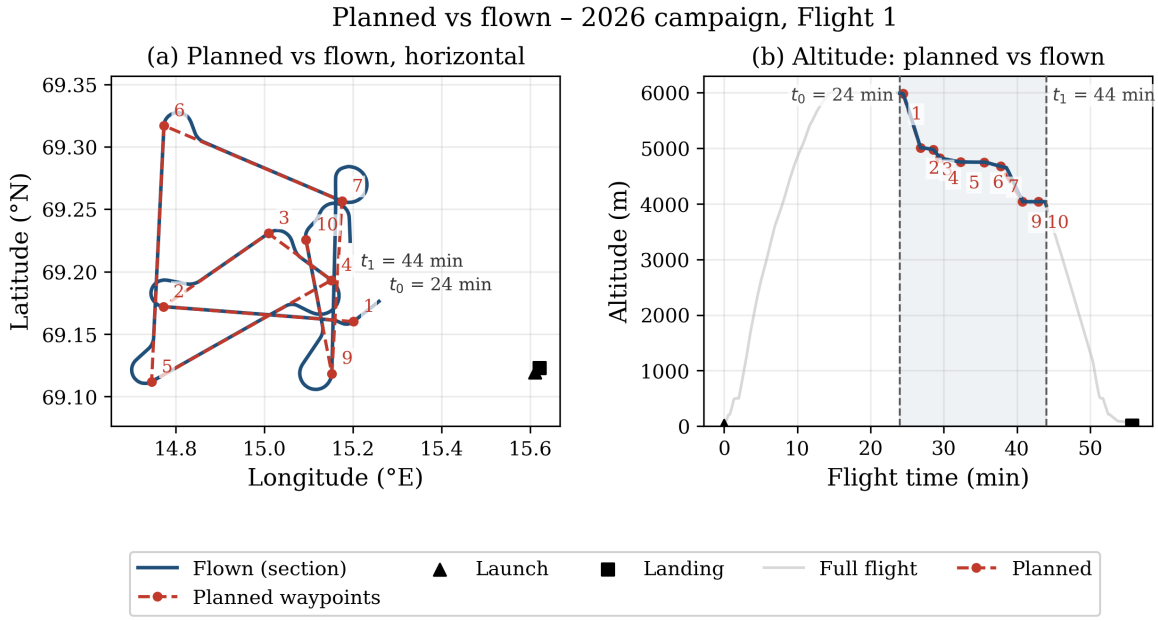


Figure 6 – Planned versus flown waypoints, 2025 Andøya campaign, Flight 1, over the section containing the optimized waypoints.

The merge of the Flight 2 measurements took 2.6 s. From measurements in hand to an upload-ready trajectory the compute was therefore about 100 s, within the 180 s minimum re-optimization interval and well under the few-minute target for an in-flight update. Measurements arrived in batches of about 60 s over the data link, and the incoming directory was polled every 30 s. One RPAS file, with placeholder slots for up to eight dropsondes, is about 25 kB, so the transfer is small and was not the limiting factor, although it was not timed separately.

The sensor profile explains the effect of the merge. Figure 9 shows the measured relative humidity against altitude. It starts near 70 % close to the surface and falls below 10 % within about 2500 m, staying low aloft. Because the surrogate field was scaled to the range $[10, 60]$, these low high-altitude values lie below the surrogate and pull the working field down where the measurements were taken. Figure 10 shows this as a north-south cross-section of the forecast and the working field, taken through the center of the mission area along the east-west axis. A clear reduction of the merged values is visible at high altitude, near 6000 m around 69.30°N, exactly where the dry measurements were introduced, as displayed in Figure 7. The merge therefore modified the surrogate in the sampled region, as intended.

The influence of the two blending scales was examined on the same data. Figure 11 compares the working field for four combinations of the horizontal and vertical scales, (σ_h, σ_v) equal to $(5\,\mathrm{km}, 200\,\mathrm{m})$, $(10\,\mathrm{km}, 200\,\mathrm{m})$, the campaign default $(10\,\mathrm{km}, 500\,\mathrm{m})$, and $(20\,\mathrm{km}, 500\,\mathrm{m})$. The smallest combination spreads the measurements too little to be visible. Increasing either scale widens the region over which the measurements modify the working field. The default sits between these extremes, local enough to follow the measurement but wide enough to leave a clear imprint.

Together the two flights confirm the operational chain. The optimizer produces flyable trajectories that the platform executes, measurements return over the data link, the merge updates the field in flight in the way the sensor profile predicts, and a re-optimized trajectory is computed, uploaded and accepted. The re-optimized route itself is not shown, as it resembles the Flight 1 waypoint list. It differs mainly in its start point, the last known RPAS position at the time of re-optimization, which can lag by up to one minute because of the upload interval.

6. Discussion

The results support three key findings. First, the system demonstrates an operational in-flight re-optimization. A trajectory was re-computed from in-flight measurements, exported, uploaded and accepted by the flight operator, which is a capability rarely shown outside simulation. Second, the

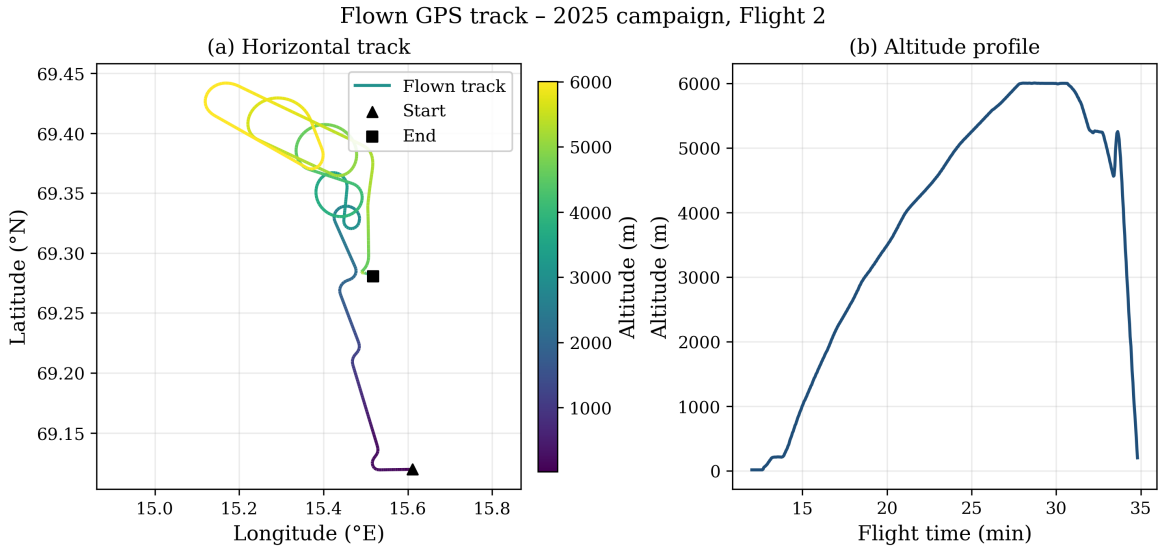


Figure 7 – GPS track of the 2025 Andøya campaign, Flight 2: horizontal track colored by altitude (left) and altitude-over-time profile (right).

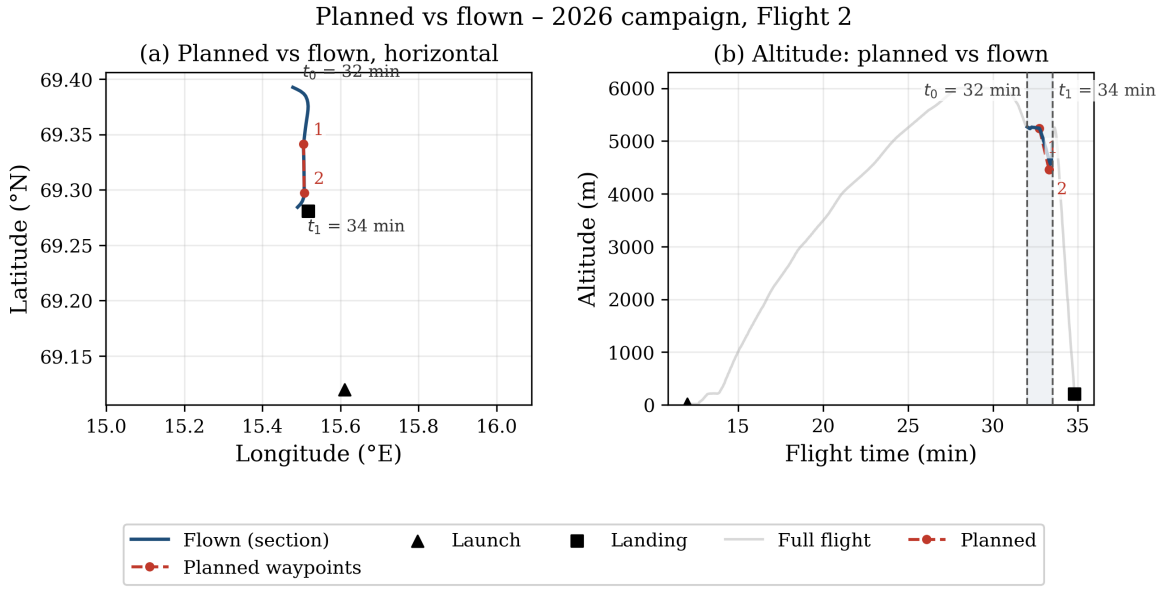


Figure 8 – Planned versus flown waypoints and altitude-over-time profile, 2025 Andøya campaign, Flight 2. Two waypoints were reached between minutes 32 and 34.

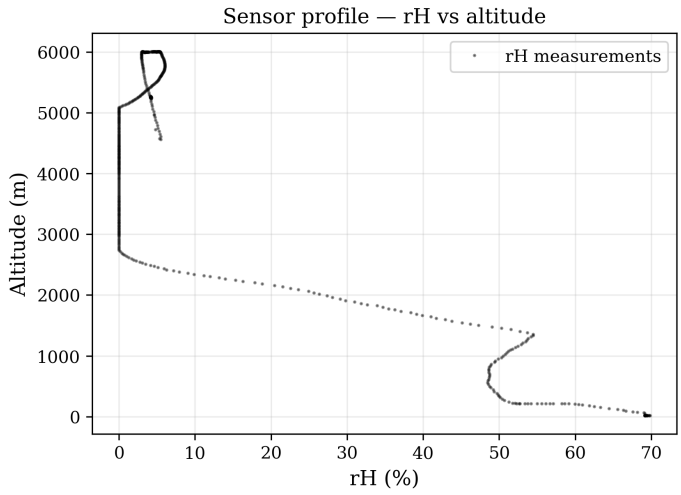


Figure 9 – Sensor profile, Flight 2: measured relative humidity against altitude. Values fall from about 70 % near the surface to below 10 % within roughly 2500 m.

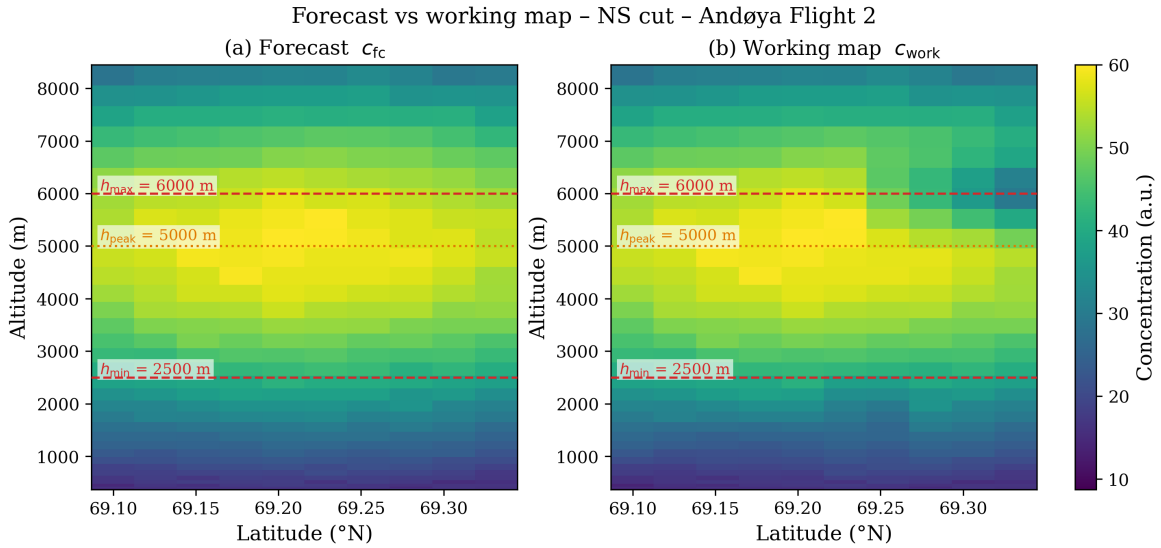


Figure 10 – North-south cross-section of the forecast field and the working field C_{work} after the first merge, Flight 2. The merged value is reduced at high altitude near 6000 m around 69.30°N, where the dry relative-humidity measurements were taken.

full chain ran end to end, from the gridded input field through three-dimensional optimization to a flight-ready, flight operator-accepted route, and back through the measurement merge. Third, the work provides an end-to-end three-dimensional methodology, integrating spatial representation, estimation, optimization, flight-dynamic constraints and navigation semantics.

A structural limit of the exploitative objective shapes the next step. $\mathrm{MAE}_{\mathrm{fit}}$ is near-constant on a spatially homogeneous field, so the objective then gives no optimization direction. This is a property of the objective, not a defect of the implementation, and it follows directly from Eq. 1. It motivates the exploration term of Section 4.6. Whether it matters in practice depends on the field. A contamination plume is strongly peaked and well suited to exploitation, whereas a smooth meteorological surrogate may not be.

Several limitations are noted factually. In the 2025 campaign the input field was a synthetic surrogate rather than a real forecast, since the operational ICON-ART interface had not yet been delivered, so the data chain is demonstrated with a surrogate standing in for the numerical weather prediction input. Flight 1 carried no GPS on the measurement computer, so positions were reconstructed from the platform GPS, which is adequate for display but not ideal for a tightly time-aligned analysis, and it is the reason no merge or re-optimization was run on Flight 1. Dropsonde profiles were not used in

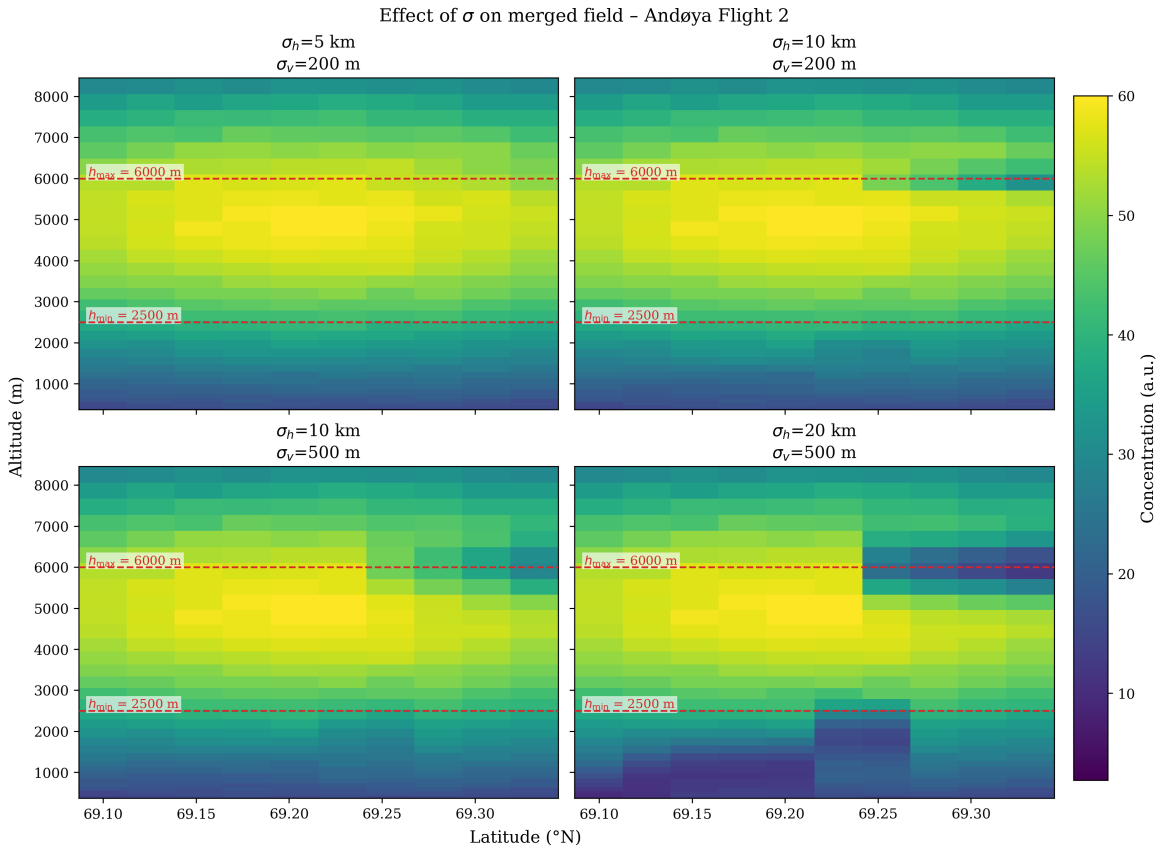


Figure 11 – Effect of the blending scales on the merged field, Flight 2, north-south cross-section. From left to right, $(\sigma_h, \sigma_v) = (5\,\mathrm{km}, 200\,\mathrm{m})$, $(10\,\mathrm{km}, 200\,\mathrm{m})$, $(10\,\mathrm{km}, 500\,\mathrm{m})$ (campaign default), and $(20\,\mathrm{km}, 500\,\mathrm{m})$.

the merge during this campaign.
Finally, the estimator choice is not yet benchmarked. Linear interpolation was selected for the reasons given in Section 4.2, but a quantitative comparison against other fast interpolation methods remains open.

7. Conclusion and Outlook

This work extends our previous two-dimensional framework to a three-dimensional, real-time adaptive trajectory optimization for RPAS-based atmospheric measurements. It integrates altitude-dependent spatial optimization, physics-based flight constraints derived from platform performance data, and navigation-aware waypoint generation. The 2025 Andøya campaign demonstrated operational feasibility: optimized trajectories were flown, measurements were merged in flight, and a re-optimized trajectory was computed, uploaded and accepted by the platform flight operator.
Future work proceeds along three lines. First, an extended campaign replaces the synthetic surrogate with real numerical weather prediction data. ICON-D2 fields serve as the input and, in a twin experiment, as the truth, so the measurement-informed analysis can be evaluated against a known field. The comparison quantifies whether the measurements reduce the error relative to the forecast, as $\mathrm{MAE}_{\mathrm{val}}$ against $\mathrm{MAE}_{\mathrm{base}}$ defined in Section 5..Operational ICON-ART composition fields are the intended input for the final operational use case. Early experience on real fields, in particular near-homogeneous temperature fields, indicates that a purely exploitative objective is not informative there, which motivates the second line.
Second, we will implement the exploration term of Eq. 5 through a low-cost visit-count surrogate, and run several adaptive update cycles within a single mission. Dropsondes have a central role in this context. Each drop yields a fast vertical profile, and several can be released per flight, so together they cover the vertical dimension that the largely horizontal RPAS track under-samples. The exploration term can target exactly that under-sampled vertical.

Third, a quantitative comparison of linear interpolation against other fast interpolation methods, with reproducible processing-time metrics. Together these steps advance the system toward operational deployment for aviation safety during atmospheric contamination events.

8. Acknowledgements

The authors thank the German Weather Service for leading the MEASURE project and for providing the meteorological forecasts and contamination scenarios, and enviscope GmbH for the construction, integration and operation of the measurement payload during the flight campaign. We further thank Airbus Defence and Space for operating the DO-TP25 platform. This work was funded by the German Federal Aeronautical Research Programme (LuFo VI-1, FKZ 20V1906B).

9. Contact Author Email Address

Kuno.Buchtal@dlr.de ; Alexander.Lau@dlr.de

10. Copyright Statement

The authors confirm that they, and/or their company or organization, hold copyright on all of the original material included in this paper. The authors also confirm that they have obtained permission, from the copyright holder of any third party material included in this paper, to publish it as part of their paper. The authors confirm that they give permission, or have obtained permission from the copyright holder of this paper, for the publication and distribution of this paper as part of the ICAS proceedings or as individual off-prints from the proceedings.

References

[1] T. Bolić and Ž. Sivčev. Eruption of eyjafjallajökull in iceland: Experience of european air traffic management. *Transportation Research Record*, 2214(1):136–143, 2011.

[2] M. Hirtl et al. A volcanic-hazard demonstration exercise to assess and mitigate the impacts of volcanic ash clouds on civil and military aviation. *Natural Hazards and Earth System Sciences*, 20(6):1719–1739, 2020.

[3] J. Schröter et al. Icon-art 2.1: a flexible tracer framework and its application for composition studies in numerical weather forecasting and climate simulations. *Geoscientific Model Development*, 11(10):4043–4068, 2018.

[4] Kuno Arthur Buchtal, Alexander Lau, Dörthe Ebert, Jochen Richters, and Kilian Schneiders. Route planning for volcanic ash and radioactivity in-situ drone measurements using a genetic algorithm and kriging. In *Proc 34th Congress of the International Council of the Aeronautical Sciences*, Florence, Italy, 2024. ISSN 2958-4647.

[5] P Tokekar, JV Hook, D Mulla, and V Isler. Sensor planning for a symbiotic uav and ugv system for precision agriculture. *IEEE Transactions on Robotics*, 32(6):1498–1511, 2016.

[6] Sebastian M Ferrandez, Tim Harbison, Thomas Weber, Robert Sturges, and R Rich. Optimization of a truck-drone in tandem delivery network using k-means and genetic algorithm. *Journal of Industrial Engineering and Management*, 9(2):374–388, 2016.

[7] K Scott, R Dai, and M Kumar. Occlusion-aware coverage for efficient visual sensing in unmanned aerial vehicle networks. *IEEE Global Communications Conference*, pages 1–7, 2016.

[8] JYJ Chow. Dynamic uav-based traffic monitoring under uncertainty as a stochastic arc-inventory routing policy. *International Journal of Transportation Science and Technology*, 5(3):167–185, 2016.

[9] Y. Zhao, Z. Zheng, and Y. Liu. Survey on computational-intelligence-based uav path planning. *Knowledge-Based Systems*, 158:54–64, 2018.

[10] Daniela Rojas Viloria, Elyn L. Solano-Charris, Andres Muñoz-Villamizar, and Jairo R. Montoya-Torres. Unmanned aerial vehicles/drones in vehicle routing problems: a literature review. *International Transactions in Operational Research*, 28(4):1626–1657, 2021.

[11] JA Guerrero and Y Bestaoui. Uav path planning for structure inspection in windy environments. *Journal of Intelligent and Robotic Systems*, 69(1–4):297–311, 2013.

[12] SH Teo, J Wang, L Geng, JYH Fuh, and YF Zhang. A heuristic mission planning algorithm for heterogeneous tasks with heterogeneous uavs. *Unmanned Systems*, 3(3):205–219, 2015.

[13] Niranjan Srinivas, Andreas Krause, Sham M. Kakade, and Matthias Seeger. Gaussian process optimization in the bandit setting: No regret and experimental design. In *Proceedings of the 27th International Conference on Machine Learning (ICML)*, pages 1015–1022, 2010.

[14] Gregory Hitz, Enric Galceran, Marie-Éve Garneau, François Pomerleau, and Roland Siegwart. Adaptive continuous-space informative path planning for online environmental monitoring. *Journal of Field Robotics*, 34(8):1427–1449, 2017.

[15] Lukas Schmid, Mihai Pantic, Raghav Khanna, Lionel Ott, Roland Siegwart, and Juan Nieto. An efficient sampling-based method for online informative path planning in unknown environments. *IEEE Robotics and Automation Letters*, 5(2):1500–1507, 2020.

[16] Eugenia Kalnay. *Atmospheric Modeling, Data Assimilation and Predictability*. Cambridge University Press, 2003.

[17] George P. Cressman. An operational objective analysis system. *Monthly Weather Review*, 87(10):367–374, 1959.

[18] Stanley L. Barnes. A technique for maximizing details in numerical weather map analysis. *Journal of Applied Meteorology*, 3(4):396–409, 1964.