Frontiers of Information Technology & Electronic Engineering
www.jzus.zju.edu.cn; engineering.cae.cn; www.springerlink.com
ISSN 2095-9184 (print); ISSN 2095-9230 (online)
E-mail: jzus@zju.edu.cn

Eixão-UAM: LLM-assisted iterative design of a low-altitude urban air mobility corridor in Brasilia*#

Li WEIGANG†‡[1], Juliano Adorno MAIA[1], Emilia STENZEL[2], Lucas Ramson SIEFERT[1]

[1]*TransLab, Department of Computer Science, University of Brasilia, Brasilia 70919-900, Brazil*

[2]*Faculty of Information Science, University of Brasilia, Brasilia 70919-900, Brazil*

†E-mail: weigang@unb.br

Received July 30, 2025; Revision accepted Oct. 20, 2025; Crosschecked Nov. 10, 2025; Published online Dec. 1, 2025

Abstract: The development of urban air mobility (UAM) systems requires scalable, regulation-aware planning of low-altitude airspace and supporting infrastructure. This study proposes an end-to-end framework for the design, simulation, and iterative optimization of a structured UAM corridor over Brasilia's central road axis (Eixão-UAM), aligned with the Brazilian unmanned aircraft traffic management (BR-UTM) ecosystem. In addition, this study proposes a multilayered aerial configuration stratified by unmanned aerial vehicle class, supported by a modular ground infrastructure composed of vertihubs, vertiports, and vertistops. A takeoff-scheduling simulator is developed to evaluate platform allocation strategies under realistic traffic and weather conditions. Initial experiments compare a round-robin (RR) baseline with a genetic algorithm (GA), and results reveal that RR outperforms GA v1 in terms of the average waiting time. To address this gap, a large language model (LLM) assisted optimization loop is implemented using GPT-4o Mini and Gemini 2.5 Pro. The LLMs act as reasoning partners, supporting the root-cause diagnoses, fitness function redesign, and rapid prototyping of five GA variants. Among these, GA v5 achieves a 59.62% reduction in maximum waiting time and an approximately 10% reduction in average waiting time over GA v1, thereby approaching the robustness of RR. In contrast, GA v2–v4 and GA v6 perform less consistently, showing an importance of fitness function design. These results underscore the role of an iterative, LLM-guided development in enhancing classical optimization, demonstrating that generative artificial intelligence (AI) can contribute to simulation acceleration and the cocreation of operational logic. The proposed method provides a replicable blueprint for integrating LLMs into early-stage UAM planning, offering both theoretical insights and architectural guidance for future low-altitude airspace systems.

Key words: Brasilia; Eixão; Genetic algorithm; Large language model (LLM); Unmanned aerial vehicle (UAV); Urban air mobility (UAM); UAM corridor; Unmanned aircraft traffic management (UTM)

https://doi.org/10.1631/FITEE.2500541 **CLC number:** TP31

1 Introduction

Urban air mobility (UAM) is rapidly emerging as a transformative component of smart urban infrastructure, enabling the use of unmanned aerial vehicles (UAVs) for logistics, emergency response, inspections, and passenger transport (Cohen et al., 2021; Jiang et al., 2023; Liao et al., 2024; Liu S and Liu, 2025; Xu et al., 2025). As cities grow denser and ground infrastructure becomes increasingly saturated, low-altitude air corridors offer a promising

‡ Corresponding author

* Project supported by the Brazilian National Council for Scientific and Technological Development (CNPq) (No. 309545/2021-8)

Electronic supplementary materials: The online version of this article (https://doi.org/10.1631/FITEE.2500541) contains supplementary materials, which are available to authorized users

ORCID: Li WEIGANG, https://orcid.org/0000-0003-1826-1850; Juliano Adorno MAIA, https://orcid.org/0000-0002-0575-7155; Emilia STENZEL, https://orcid.org/0009-0009-5370-2595; Lucas Ramson SIEFERT, https://orcid.org/0009-0009-8038-9506

© Zhejiang University Press 2025

solution to mitigate surface congestion and improve service responsiveness (Verma et al., 2022; Moon et al., 2024). However, implementing scalable and safe UAM systems poses significant challenges in terms of airspace structuring, flight intent management, and ground infrastructure deployment.

Brazil is uniquely positioned to lead UAM integration through its national unmanned aircraft traffic management (UTM) framework, known as BR-UTM (da Silva et al., 2020; Jasper and Nunes, 2022; Turchetti and Murça, 2024). This framework features key platforms such as unmanned aircraft system for UAV registration and authorization (SISANT), which is managed by the National Civil Aviation Agency of Brazil (ANAC), and the request for access for remotely piloted aircraft-next generation (SARPAS NG), which is operated by the Brazilian Department of Airspace Control (DECEA). Both the ANAC and DECEA are interoperable with the national Gov.br digital infrastructure. Leveraging this regulatory and technological foundation, Brasilia, the capital of Brazil and a city designed with geometric precision, offers an ideal testing ground for structured low-altitude airspace models.

To explore these opportunities, this study concentrates on the modeling and validation of a UAM corridor along Brasilia's iconic south–north road axis (Eixão Rodoviário, also known as Eixão, Portuguese for "big axis"). The primary objective is to design and assess the BSB Eixão-UAM corridor within Brazil's BR-UTM framework, ensuring alignment with national airspace management initiatives.

While the corridor design remains the central focus, the analysis of takeoff scheduling at a vertiport is presented as the first operational demonstration. This addition grounds the corridor concept in practical feasibility, demonstrating how airspace architecture and ground infrastructure interact under realistic traffic scenarios. For consistency with international spelling convention, the corridor is hereafter referred to as Eixão-UAM throughout the paper.

The proposed model integrates a vertically and laterally stratified airspace with a modular ground infrastructure composed of vertihubs, vertiports, and vertistops. The Eixão-UAM simulation framework, which assesses traffic flow, takeoff scheduling, and platform utilization under both normal and adverse conditions, is developed to evaluate system performance (Fig. 1).

Methodological overview is shown as follows. Beyond conventional UAM architectural designs, a scheduling analysis method is incorporated to validate corridor feasibility. Two approaches are benchmarked: a round-robin (RR) baseline (Liu CL and Layland, 1973; Halder et al., 2023; Varun Kumar et al., 2023) and a genetic algorithm (GA) optimizer (Ferreira et al., 2014; Alolaiwy et al., 2023; Chan et al., 2023). In simulations with 5000 daily UAV departures, the RR consistently achieves lower average waiting time than the GA, highlighting a misalignment in the fitness function of the GA.

To address this, the large language models (LLMs, such as GPT-4.5, Gemini 2.5 Pro, and GPT-4o Mini) are employed to iteratively redesign GA variants. While lightweight models emphasize throughput at the expense of fairness, larger models capture multi-objective trade-offs more effectively. This process underscores the value of LLMs as code generators and reasoning partners for algorithmic exploration and refinement.

Contributions of this work are listed as follows. This study introduces the Eixão-UAM framework, integrating architecture-aware corridor planning, simulation-based scheduling, and LLM-assisted optimization. The main contributions of this work are threefold. (1) This study demonstrates that the quality of the cost function, its alignment with system-level objectives, has a greater impact on scheduling performance than algorithmic complexity. (2) Through an iterative LLM-guided design, GA v5 achieves a 59.62% reduction in maximum waiting time and an approximately 10% reduction in average waiting time compared with GA v1, narrowing the gap with the robustness of RR. (3) The study demonstrates that LLMs can serve as code generators and conceptual partners, enabling human–AI (AI is short for artificial intelligence) collaboration in exploring trade-offs and improving fairness in the design of UAM systems.

2 Related works

The evolution of UAM has progressed from institutional frameworks to air traffic management (ATM) extensions, and now toward innovations for low-altitude urban operations. This section reviews four dimensions: (1) international regulatory efforts,

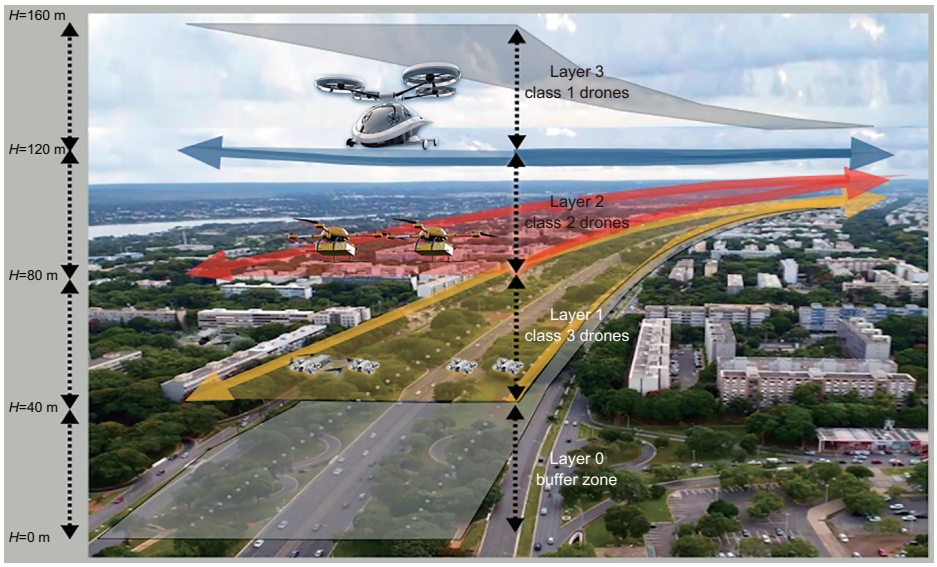


Fig. 1 Concept of the Brasilia Eixão UAM corridor (Eixão-UAM), an integrated urban ground-air mobility corridor leveraging the central highway axis (Eixão). General layout showing the location of both aerial corridors along the green strips. References to color refer to the online version of this figure

(2) Brazil's national UTM integration, (3) computational methods in UAM operations, and (4) recent advances in LLM-assisted UAM research.

2.1 Global regulatory frameworks in UAM

Global regulatory bodies have laid important groundwork for UAM integration. The manual on remotely piloted aircraft systems, Doc 10019, written by the International Civil Aviation Organization (ICAO), provides the foundational principles for harmonizing UTM across regions (ICAO, 2015). In addition, the Federal Aviation Administration (FAA) in the United States published its phased UAM ConOps 2.0 and implementation plan, while NASA has advanced operational testing through the UAM grand challenge and the UAM maturity level framework (Lascara et al., 2018; Goodrich and Theodore, 2021; FAA, 2023a, 2023b).

Europe's SESAR program has introduced the U-space architecture (U1–U4) through defining service layers for automation and real-time traffic management (EASA, 2017). Meanwhile, China has also demonstrated operational maturity in UAM deployments, notably in cities such as Shenzhen, where air corridors have been integrated into a broader urban mobility strategy (Cohen et al., 2021; Yang et al., 2024). These early deployments have resulted in logistical gains and social benefits, such as improved access to markets for rural producers.

2.2 Brazilian BR-UTM architecture

The Brazilian UTM (BR-UTM) ecosystem serves as the national framework for integrating UAVs into both controlled and uncontrolled airspaces. It is primarily coordinated by DECEA and ANAC, in accordance with ICAO's UTM guidelines (da Silva et al., 2020; Jasper and Nunes, 2022; Turchetti and Murça, 2024). Brazil demonstrates a distinctive model of integrated UTM development, guided by a phased institutional roadmap and supported by digital governance mechanism. Since 2009, DECEA has progressively issued regulatory instruments—beginning with AIC N29/09 and culminating in ICA 100-40 (DECEA, 2023)—to establish safe UAV access to national airspace. Complementing these initiatives, ANAC enacts RBAC-E94 (ANAC, 2017b) and develops the SISANT platform, which enables registration and regulatory compliance for UAVs weighing under 25 kg.

At the core of Brazil's operational UTM implementation is SARPAS NG, a web-based platform that enables UAV operators to request flight authorizations, manage airspace deconfliction, and interface with restricted zones (de Vasconcellos and Regis, 2022). SARPAS NG is fully integrated into the federal Gov.br ecosystem, aligning with

Brazil's national digital transformation strategy and providing unified access to the following services:

(1) UAV registration and operator identification;

(2) Flight planning and approval workflows;

(3) Strategic and tactical airspace conflict resolution;

(4) Priority processing for critical missions.

2.3 Emerging computational methods in UAM operations

Recent technical literature documents a shift from rule-based ATM logic to AI-augmented, data-driven approaches for UAM planning. Notably, NASA's X4 campaign emphasizes demand-capacity balancing in congested urban airspaces (Cheng et al., 2022). Along similar lines, van Nguyen (2020) introduced dynamic delegated corridors (DDCs) and four-dimensional (4D) required navigation performance (RNP) routes to improve predictability and reduce controller workload.

In conflict management, an AI-based approach, mainly reinforcement learning (RL), is applied to adaptively and smartly integrate four types of deconflict actions for solving conflicts with fewer environmental delays (Liang et al., 2019). Deep RL (DRL) is emerging as a key enabler. Studies have applied multi-agent actor–critic models (A3C) (Deniz and Wang, 2024), energy-aware destination matching with DRL backbones (Zou et al., 2024), and reward-optimized policy networks for dense UAM navigation (Garcia et al., 2023). Decentralized solutions such as the estimated time of arrival (ETA)-based self-separation at constrained waypoints further demonstrate the potential of distributed control schemes (Pruekprasert and Nakadai, 2024).

The corridor-based airspace design is increasingly recognized as a practical strategy for structured UAM flows. Examples include FAA-aligned concepts for Dallas–Fort Worth (Verma et al., 2022), grid-based skyline architectures (Muna et al., 2021), and regional corridors considering noise and social acceptance (Korringa et al., 2025). Comparative studies have further highlighted the trade-offs among slot-based, trajectory-based, and corridor-based schemes in terms of scalability and automation potential (Abdellaoui et al., 2025).

This emerging body of work highlights the centrality of corridor-oriented and scheduling-focused methods, which directly motivate the structured design and optimization framework developed in this paper.

2.4 Advances in UAM with LLM integration

Conventional approaches rely on queue-based scheduling (Liu CL and Layland, 1973), heuristic methods such as GA and ant colony optimization (ACO) (Ferreira et al., 2014), and 4D trajectory planning with environmental constraints (Alolaiwy et al., 2023). However, the integration of LLMs into simulation-based decision pipelines remains limited.

Recent studies have explored the application of LLMs in UAV-related domains. Gong et al. (2025) proposed a hybrid DRL-LLM framework for generating safe UAV trajectories under regulatory constraints. Sadik et al. (2025) introduced a holonic UAM architecture that employs LLMs as deliberative agents for cooperative task allocation. Moraga et al. (2025) used LLMs to assist in traffic-aware UAV dispatching, aiming to mitigate congestion and reduce emissions in smart city environments.

Although these works are promising, they remain largely conceptual or focus on semantic tasks such as prompt-response reasoning, rather than real-time decision loops. To date, no study has applied LLMs to vertistop and vertiport scheduling, air pad allocation, or the iterative refinement of cost functions and fitness metrics in UAV launch sequencing under dynamic constraints.

To the best of our knowledge, this study is the first to integrate LLMs within an iterative, simulation-driven optimization framework for vertiport-based UAV scheduling. Unlike conventional uses of LLMs as coding assistants or static prompt generators, the proposed architecture employs LLMs as semantic partners in engineering design, supporting root-cause diagnoses, cost function tuning, hyperparameter adaptation, and statistical performance evaluations.

3 Architectural layout and corridor planning

3.1 Brasilia Eixão: relevant information

The Eixão, designed by Lúcio Costa in 1957, is the central north–south axis of Brasilia and a key component of the city's UNESCO World

Heritage-protected urban plan (Fig. 1). Extending 13.8 km, it links the residential and institutional sectors of the north and south wings, complementing the Monumental Axis by serving as the backbone of daily urban mobility.

Morphologically, the Eixão constitutes a high-capacity linear road system comprising a seven-lane expressway (80 km/h), flanked by 45-m wide green buffers, and auxiliary four-lane service roads known as Eixinhos (60 km/h). A series of overpasses (trevos) connect the expressway to local superblocks, while metro stations, public transport corridors, and pedestrian crossings facilitate multimodal integration. On weekends, the expressway is closed to vehicular traffic, transforming into Brazil's largest open-air recreational corridor.

This distinctive combination of linearity, centrality, and unobstructed geometry renders the Eixão particularly suitable for structured UAM corridors. Its spatial clarity and functional density reflect the modernist principles underpinning Brasilia's urban design while providing a robust foundation for low-altitude aerial integration and future mobility innovations.

3.2 Overview of the Eixão-UAM corridor

The proposed aerial corridor over the Eixão is developed through an integration of technical, spatial, regulatory, and urban planning criteria, aligned with international UAM guidelines and adapted to the heritage-protected pilot plan of Brasilia (Fig. 2). These parameters ensure physical feasibility, operational safety, and institutional compliance within a consolidated urban environment.

The Eixão's linear morphology and central green buffers form a natural longitudinal channel for low-altitude electric vertical takeoff and landing (eVTOL) operations, minimizing turbulence and ground interference. Its spatial clarity and functional land-use segregation promote predictable traffic patterns and reduce potential conflicts with existing urban functions, while surrounding low-density and institutional zones enhance operational safety margins.

Strategically positioned vertiports and intermediate stops can be aligned with key demand hubs, such as government centers, hospitals, and transit terminals, to ensure functional coherence and multimodal connectivity. Existing infrastructure, including metro stations, bus rapid transit (BRT) corridors, and pedestrian underpasses, provides opportunities for seamless last-mile integration and inclusive urban access.

Owing to its centrality and symbolic status, the Eixão corridor is also suitable for phased pilot projects involving regulatory and planning agencies. A stepwise implementation strategy can facilitate validation, governance experimentation, and gradual social acceptance, positioning the Eixão as a reference model for UAM in Latin America.

3.3 Airspace and infrastructure composition

The proposed aerial corridor is designed around the Eixão's linear geometry, which supports the implementation of parallel, unidirectional airways along the central green strips. The western strip is designated for the north-to-south traffic, and the eastern strip for the south-to-north traffic.

The decision to represent the flight corridors with equal spacing is grounded in three complementary aspects: (1) UAV operational safety, (2) regulatory alignment and scalability, and (3) structural efficiency and urban design coherence. Following the tiered approach outlined in DECEA's ICA 100-40 and consistent with international UAM models, the proposed segmentation (0–40, 40–80, 80–120, and 120–160 m) ensures functional separation between UAV classes. The lowest band (0–40 m) acts as a buffer for takeoff, landing, and emergency maneuvers, whereas the subsequent layers, from second lowest to highest, are allocated to lightweight, medium, and heavy UAVs, respectively, reflecting their different payloads, velocities, and safety requirements. Although DECEA has not yet formalized class-specific airspace separations, this proposal anticipates forthcoming regulatory developments and establishes a scalable framework for future urban operations. Equal spacing further simplifies conflict detection, minimizes overlapping flight envelopes, and enhances the reproducibility of simulation scenarios under varying traffic conditions (Roberge et al., 2013; Lavezzi et al., 2023).

Beyond safety and regulation considerations, equal spacing also offers infrastructural and design advantages. A grid-based vertical layout enables modular expansion, reduces the need for specialized structural adjustments, and facilitates integration with vertiports and navigation systems. From an architectural standpoint, regularly spaced elements

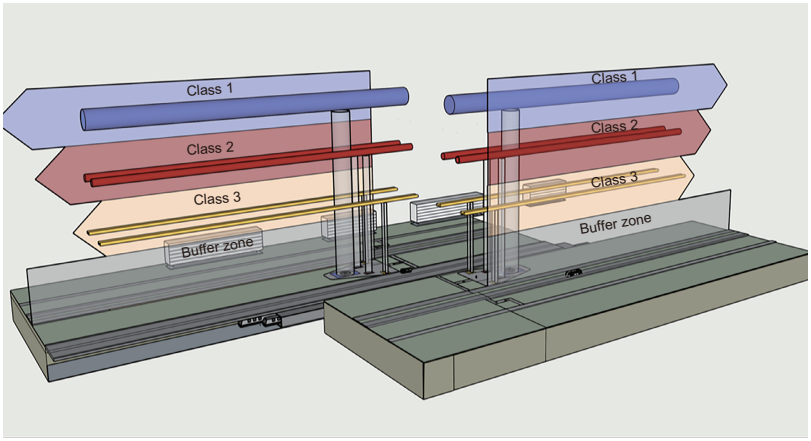


Fig. 2 Three-dimensional (3D) diagram showing ascent and descent flows and a layered traffic structure

introduce predictability, visual clarity, and adaptability, thereby reinforcing both operational efficiency and public acceptance of emerging aerial infrastructure (Neufert, 2013). Overall, the proposed corridor design balances engineering rigor, regulatory foresight, and spatial coherence.

Each corridor is vertically stratified into four layers of approximately 40 m each, arranged in a stacked configuration to accommodate different classes of UAVs in accordance with ANAC's regulatory standards (ANAC, 2017b). The detailed classification criteria are presented in Section 4. This vertical segmentation is consistent with DECEA's national regulatory framework (DECEA, 2023), ensuring alignment with the BR-UTM ecosystem and supporting safe, predictable traffic flows by differentiating operations according to UAV size and function (Fig. 2).

1. Layer 0 (0–40 m): not used for cruise flights. This layer serves as a vertical access and transition zone, supporting takeoff and landing procedures while buffering noise and visual impacts.

2. Layer 1 (40–80 m): reserved for class 3 UAVs (light drones). This layer is used in last-mile deliveries, public services, and high-maneuverability tasks, and offers maximum operational flexibility.

3. Layer 2 (80–120 m): allocated to class 2 UAVs (medium drones). This layer supports mid-range routes with higher payloads and more predictable paths.

4. Layer 3 (120–160 m): reserved for class 1 UAVs (heavy drones or passenger eVTOLs). This layer requires greater vertical and lateral separation, owing to their size and speed.

This vertical layering supports traffic growth and UAV specialization while reducing conflicts and streamlining ascent, cruise, and descent operations. Vertical access is provided at designated vertistops via Layer 0, ensuring controlled integration into higher layers.

Vertistops are placed in high-demand zones such as hospitals, government areas, and commercial centers, each with adjacent escape zones for emergency rerouting (Fig. 3).

As Brasilia's central axis, the Eixão corridor also functions as a visible pilot site, with a modular design that allows phased activation by drone class, area, or time, facilitating gradual regulation, stakeholder coordination, and public acceptance of UAM systems.

3.4 Vertihub, vertiport, and vertistop designs

The distribution of vertiports and vertistops along the BSB Eixão-UAM corridor is guided by both Brasilia's urban morphology and functional mobility demand. Conceived by Lúcio Costa and recognized as a UNESCO World Heritage site, the Plano Piloto features a linear north–south axis (Eixão) flanked by 45-m wide green buffers and modular superblocks. These characteristics naturally support a layered UAM design: Intersections of commercial streets and superblocks yield 15–16 potential vertistops, whereas wider buffer zones accommodate modular pads with minimal acoustic and visual impact. Table 1 presents the spatial coordinates of the UAV vertihub, vertiports and vertistops along

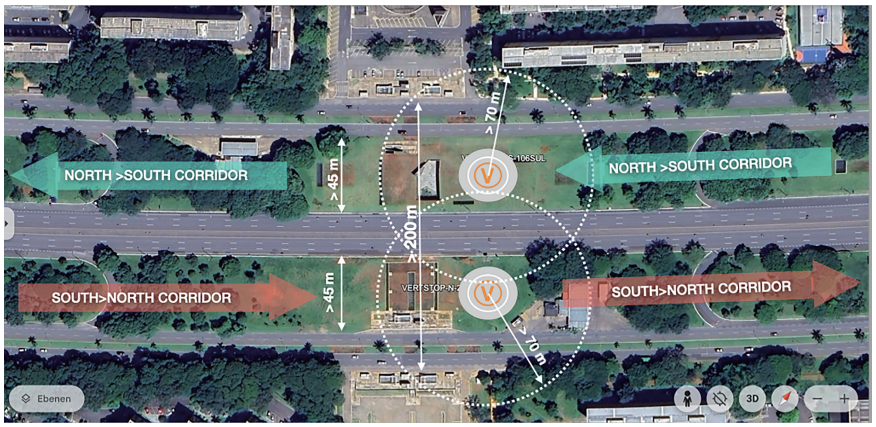


Fig. 3 Proposed distribution of UAM infrastructure along the Eixão corridor: 15 modular vertistops embedded in the green buffer strips, three vertiports (north, central, and south) integrated with intermodal hubs, and one vertihub at the road–rail terminal. The configuration balances functional demand, multimodal connectivity, and spatial coherence

Table 1 Spatial coordinates of the UAV vertihub, vertiports, and vertistops in Eixão-UAM

Point type	Direction	Location	Latitude	Longitude
South wing vertiport		Interchange EPIA/EPGU	15°50′26″S	47°55′45″W
Vertistop - 102 south	N → S	South axis (Eixão)-SQS 102	15°48′21″S	47°53′22″W
Vertistop - 202 south	S → N	South axis (Eixão)-SQS 202	15°48′23″S	47°53′21″W
		...		
Central vertihub		Bus and railway terminal	15°46′31″S	47°56′31″W
Central vertiport		Plano Piloto bus terminal	15°47′39″S	47°52′52″W
		...		
Vertistop - 110 north	N → S	North axis (Eixão)-SQN 110	15°45′19″S	47°53′11″W
Vertistop - 210 north	S → N	North axis (Eixão)-SQN 210	15°45′18″S	47°53′09″W
		...		
North wing vertiport		Roundabout TORTO/EPIT	15°43′58″S	47°53′40″W

EPIA, EPGU, TORTO, and EPIT are Brasilia highways. N, W, and S denote the north, west, and south, respectively

the proposed Eixão-UAM corridor.

Complementing these distributed nodes, three vertiports and one vertihub are defined by criteria of intermodality, activity density, and metropolitan connectivity. The northern and southern vertiports are positioned near major junctions linking the pilot plan (Plano Piloto—center of Brasilia city) with satellite cities, adjacent to mixed-use complexes of shopping centers, healthcare facilities, and high vehicular flows. The central vertiport is co-located with Brasilia's main intermodal bus and metro terminal, directly serving commercial and government districts. The proposed vertihub is situated next to the Rodoferroviária (road–rail terminal), connecting logistics infrastructure with industrial sectors and regional road corridors such as DF-095, thereby extending system reach beyond the central area.

This preliminary blueprint illustrates how Brasilia's spatial logic-linearity, modularity, and green buffers can be adapted to UAM infrastructure. The final implementation, however, must undergo regulatory, environmental, and socio-technical evaluations under government coordination to ensure safe, scalable, and publicly accepted deployment.

3.5 Broader considerations and strategic implications

The Eixão UAM corridor represents a landmark initiative in Latin America, leveraging Brasilia's linear morphology, zoning logic, and multimodal connectivity to enable scalable UAM deployment. The design incorporates vertical stratification by UAV class, green buffer integration, and strategically positioned takeoff points, aligning with international best practices while preserving the city's World Heritage integrity.

As a pilot project, it promotes coordination among federal regulators, local authorities, and

community stakeholders, positioning Brasilia as a modernist urban laboratory for three-dimensional (3D) mobility. Future research directions include noise mitigation through buffer zoning, renewable-energy integration for vertiport charging, and UAV applications in emergency logistics.

4 Methodology

This section details the methodology developed to simulate UAV flight operations and scheduling across a network of distributed vertiports along the Brasilia south–north axis (Maia et al., 2025).

4.1 Flight plan structure and spatial constraints

To simulate conflict-free operations, each UAV flight plan must specify critical parameters, including geographic displacement, vehicle class, flight altitude, and departure timing. UAV movement is assumed to occur along a single axis (either latitude or longitude), and velocity is defined as

$$\boldsymbol{V} = \frac{\boldsymbol{L}_{\mathrm{a}} \times \boldsymbol{L}_{\mathrm{o}}}{t}, \tag{1}$$

where $\boldsymbol{L}_{\mathrm{a}}$ and $\boldsymbol{L}_{\mathrm{o}}$ denote displacement vectors in the latitude and longitude, respectively, and t is the travel time.

UAV velocity is bounded by class-specific operational constraints. Table 2 summarizes Brazil's UAV classification by maximum takeoff weight and corresponding regulatory altitude bands (ANAC, 2017a). Vertical separation is assumed to be at least equal to the UAV height, based on safety margins in the absence of mandatory minimum vertical separation regulations.

A vertically layered airspace model is adopted based on the geometry of the Eixão corridor (approximately 13.8-km long and 250-m wide), allowing opposing traffic flows to be separated by the altitude rather than the lateral position. In this 3D configuration, each UAV class is assigned to a dedicated vertical lane, ensuring conflict-free bidirectional movement. Vertiports are located at the two corridor terminals and at the midpoint of the Eixão. Intermediate vertistops are placed every two superblocks to support routine takeoff and landing operations and to provide distributed access along the corridor. The detailed spatial distribution is summarized in Table 1.

Table 2 UAV classification by weight and operational altitude

Class	Max takeoff weight	Min altitude (m)	Max altitude (m)
1	Greater than 150 kg	120	140
2	From 25 up to 150 kg	80	119
3	From 0.25 up to 25 kg	40	79

Source: ANAC (2017a)

4.2 Vertiport configuration and takeoff scheduling

To reflect capacity limitations at each vertihub and vertiport, the number of available takeoff pads is defined based on the UAV class, as summarized in Table 3.

Takeoff requests are prioritized based on the following hierarchical rule set presented in Table 4, consistent with the logic adopted by Brazil's SARPAS system.

Two scheduling algorithms are evaluated to manage takeoff slot allocation under variable demand: an RR method and a GA-based optimizer.

4.3 Round-robin scheduling mechanism

The RR algorithm cycles through all takeoff requests in the priority queue, assigning each a fixed time slice (quantum) of 30 s. If a UAV cannot complete its launch in one quantum, it is re-queued with its remaining required time updated.

Formally, for each request P_i during scheduling cycle c, the service duration for request P_i is defined as follows:

$$P_i = \min\left(q,\ t_i^{(c)}\right), \tag{2}$$

where $q = 30\,\mathrm{s}$ and $t_i^{(c)}$ is the remaining time in cycle

Table 3 Number of launch pads per vertiport, by UAV class

Class	Number of launch pads
1	1
2	2
3	4

Table 4 Priority rules for takeoff slot allocation at vertiports

Priority level	Description
1	Official or emergency aircraft (e.g., public safety or health missions)
2	Heavier and larger UAVs
3	Scheduled departure time

c. If $t_i^{(c)} > q$, the remaining time is updated for the next cycle:

$$t_i^{(c+1)} = t_i^{(c)} - q. \tag{3}$$

This approach ensures fairness and prevents takeoff slot monopolization. RR methods are widely adopted in UAV coordination studies, such as dynamic super round UAV clustering (DROVE) and Emerald queueing systems (Halder et al., 2023; Varun Kumar et al., 2023).

4.4 Genetic algorithm-based optimization

The GA approach is used to determine the optimal sequence of drone launches which minimizes the total waiting time and pad occupancy penalties. The optimization problem is modeled as follows (Maia et al., 2025):

Let $Q = \{a_1, a_2, \ldots, a_{N_{\text{queue}}}\}$ denote the current queue of UAVs. S_Q refers to the set of all feasible scheduling sequences that can be formed from the current queue Q. The goal is to find a sequence c^* minimizing the objective function $f(c)$.

$$c^* = \arg\min_{c \in S_Q} f(c). \tag{4}$$

The cost function $f(c)$ is defined as

$$f(c) = \sum_{i=1}^{N_{\text{queue}}} \left(t_{\text{departure}}(a_i', c) - t_{\text{arrival}}(a_i')\right) + \sum_{k \in K} w_k \cdot \left(\max_{p \in P_k}(t_{\text{free}}) - t_{\text{current}}\right), \tag{5}$$

where N_{queue} is the number of UAVs currently waiting in the queue, $t_{\text{departure}}(a_i', c)$ denotes the scheduled departure time for UAV a_i' under sequence c, $t_{\text{arrival}}(a_i')$ denotes the request (arrival) time for UAV a_i', K is the set of vertiport facilities, P_k denotes the set of service pads associated with facility k, t_{free} is the next available (free) time of a pad in P_k, t_{current} is the present simulation time, and w_k is the penalty weight for facility k. The second term thus penalizes delays in resource reuse when the latest t_{free} among pads in P_k exceeds t_{current}.

GAs are suitable for this class of scheduling problems (Ferreira et al., 2014; Alolaiwy et al., 2023). They offer flexibility for future integration with dynamic weather, emergency prioritization, or energy-aware planning.

5 Preliminary results and discussion

This section presents and discusses the results of a 30-d simulation of UAV operations at a high-demand vertiport, comparing the performance of a standard scheduling algorithm (RR) with that of a GA designed to optimize takeoff sequencing in different conditions.

5.1 Simulation framework and justification

5.1.1 General configuration and traffic structure

To ensure realism and relevance, the simulation was configured to reflect the operational conditions of a high-capacity urban vertiport. A total of 5000 drone flights per day are simulated over a 30-d period, spanning operational hours from 06:00 to 22:00.

Fig. 4 illustrates a representative layout of launch pads within a vertiport, used for scheduling and simulation analyses. The notation system follows the pattern HP-Cx-y, where HP denotes the home point used for vertical takeoff and landing, Cx indicates the UAV class (e.g., C1 for class 1, C2 for class 2), and y is the unique identifier for each launch window within the same class. In this example, HP-C1-1 indicates a single launch pad dedicated to class 1 UAVs (typically large, passenger or heavy-cargo drones); HP-C2-1 and HP-C2-2 indicate two launch pads for class 2 UAVs (medium-range, logistics drones); HP-C3-1 through HP-C3-4 indicate four launch pads for class 3 UAVs (lightweight, short-range drones).

A global safety delay of 30 s is applied between consecutive takeoffs to model airspace separation procedures and wake turbulence management.

Table 5 presents the allocation of launch pad capacity per UAV class and the penalty weights used in the GA model.

Table 5 Takeoff parameters used in the flight simulations

Drone class	Launch pad share (%)	GA penalty weight
Class 1	10	1.5
Class 2	30	1.0
Class 3	60	0.5

Source: adapted from Ferreira et al. (2014)

The number and spatial layout of launch pads for each UAV class are treated as optimization variables in this study. Scheduling algorithms—

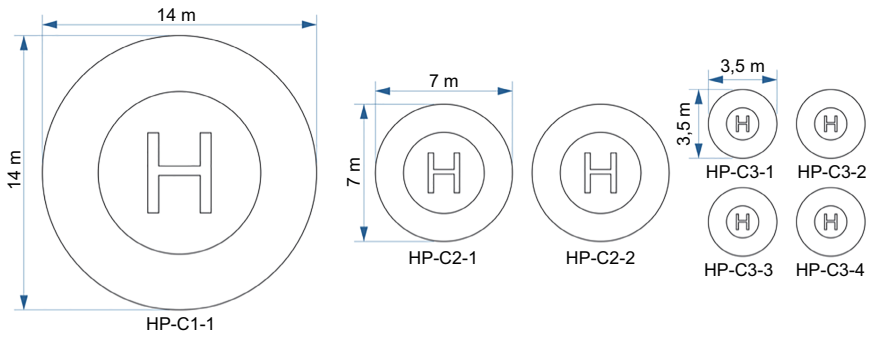


Fig. 4 Schematic layout of launch pads inside a vertiport, categorized by UAV class. Launch pads are labeled as HP-Cx-y, where x indicates UAV class and y indicates the pad ID within that class

including RR, GA, and LLM-assisted planners—are applied to determine optimal launch pad configurations that minimize waiting time and enhance overall operational efficiency.

5.1.2 Flight schedule and environmental conditions

Flight demand exhibits a bimodal peak pattern, with morning (07:00–09:00) and evening (17:00–19:00) rush hours. Weather events are generated probabilistically, following the parameters outlined in Table 6.

During the simulation period, 10 major weather events occur (Table 7), including storms on Jan. 3 and 18, and strong winds on Jan. 4 and 30.

Meteorological variability is further modeled using randomly triggered events based on relative humidity values ranging from 29.4% to 32.3%, characteristic of Brasilia's rainy season (Bayer and Bayer, 2015). These adverse weather periods result in numerous departures being delayed or canceled, reflecting realistic operational disruptions.

5.1.3 Fleet composition and randomized disruptions

The simulated fleet comprised: (1) 2% high-priority official flights (e.g., emergency or governmental missions) and (2) 10%, 30%, and 60% for classes 1, 2, and 3, respectively, reflecting the predominance of light UAV traffic in urban environment.

Table 6 Impact modeling for each type of weather event

Event type	Pads closed (%)	Delay multiplier
Heavy rain	50	1.25×
Storm	75	1.50×
Moderate wind	0	1.25×
Strong wind	0	1.50×

Table 7 Weather event log during the 30-d simulation period

Date	Event type	Time (start–end)
03/01/2025	Storm	19:51–20:52
04/01/2025	Strong wind	17:45–18:50
11/01/2025	Heavy rain	19:09–20:26
13/01/2025	Moderate wind	12:49–13:55
17/01/2025	Moderate wind	19:38–20:19
18/01/2025	Storm	10:03–11:06
19/01/2025	Moderate wind	16:29–17:37
20/01/2025	Moderate wind	15:10–16:00
25/01/2025	Moderate wind	18:45–20:14
30/01/2025	Strong wind	08:04–09:08

Weather-induced disruptions are modeled using a 25% daily probability of occurrence, consistent with climatological records from the National Institute of Meteorology (INMET) for Brasilia international airport (SBBR). This ensures that simulated weather conditions reflected seasonal patterns in the target region.

5.1.4 Justification of weather impact modeling

The modeled impacts of weather events are grounded in standard aviation safety protocols and UAV literature.

1. Rain and storms: These events reduce vertiport capacity by closing 50%–75% of pads. This reflects decreased visibility, slippery surfaces, and the need for increased spacing between takeoffs because of safety procedures.

2. Winds (moderate and strong): Wind effects are modeled by increasing separation times between UAV operations, applying delay multipliers of 1.25× for moderate winds and 1.50× for strong winds. This approach aligns with findings from Wang et al.

(2019), which indicate that crosswinds and gusts elevate flight control difficulty, demanding greater attention from both pilots and air traffic controllers. Accordingly, longer intervals are incorporated as a safety measure. The higher incidence of moderate wind events in the simulation also reflects historical climatological records for the region.

5.2 Comparative performance analysis

The RR scheduler serves as an idealized benchmark for scheduling but does not account for operational constraints such as priority classes, time-varying demand, or penalty structures. In contrast, GA-based approaches explicitly incorporate these factors, trading off isolated performance metrics for regulatory and operational realism. The simulation results (Table 8) indicate the low average delays across methods, with RR retaining an advantage in execution time for this allocation problem.

5.2.1 Average waiting time by UAV class

Fig. 5 shows the average waiting time (in minutes) per UAV class, comparing the performance of RR and GA scheduling. A quantitative breakdown reveals the following.

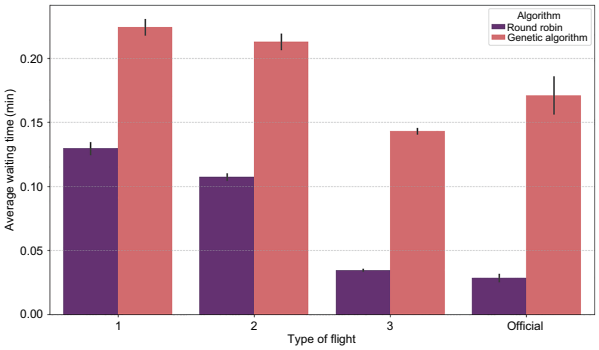


Fig. 5 Average waiting time by UAV class for RR and the GA

1. Class 1: RR outperforms the GA significantly (≈0.13 min vs. 0.22 min).

2. Class 2: RR achieves 0.11 min vs. 0.21 min for the GA.

3. Class 3: RR exhibits a significant advantage (0.03 min vs. 0.14 min).

4. Official flights: RR has an advantage (0.025 min vs. 0.17 min).

This pattern suggests a consistent advantage for RR in minimizing delay, particularly for high-priority and small-class flights.

5.2.2 Aggregated hourly throughput

Fig. 6 illustrates the total cumulative hourly throughput over 30 d for each hour of the days. The blue and orange lines represent RR and GA scheduling, respectively.

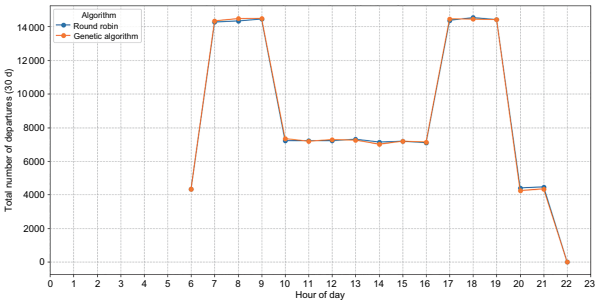


Fig. 6 Total hourly throughput over a 30-d period. References to color refer to the online version of this figure

The pattern is bimodal, which is consistent with typical urban mobility peaks.

1. 07:00–09:00: first peak with over 14 000 take-offs.

2. 10:00–16:00: midday stabilization around 7000 per hour.

3. 17:00–19:00: second, more intense peak again exceeding 14 000.

At the macro level, both algorithms exhibit nearly identical throughput, indicating a similar handling capacity. However, significant differences arise under adverse weather conditions.

5.3 Performance under different weather conditions

Fig. 7 compares throughput across different weather scenarios. Key insights include the following.

1. No event: identical performance, confirming equivalence under unrestricted conditions.

2. Moderate wind: mixed results, with RR leading in some periods and the GA in others.

3. Strong wind: GA manages peaks slightly better, but RR achieves shorter waiting times.

4. Heavy rain: RR has a slight advantage.

5. Storm: GA leads in isolated peaks (e.g., 19:00) but underperforms at other times.

Table 8 Numerical comparison panel between the two allocation algorithms

Metric	RR	GA
Overall performance		
Average waiting time (general) (s)	4	7
Median waiting time (s)	0	0
Average throughput (takeoffs/h)	313	313
Consistency (Std. dev. of waiting time) (s)	13	15
Worst case (max waiting time) (s)	131	105
Performance by category		
Average wait for official flights	2	9
Average wait for class 1 (s)	8	11
Average wait for class 2 (s)	9	8
Average wait for class 3 (s)	1	6

Std. dev. denotes standard deviation

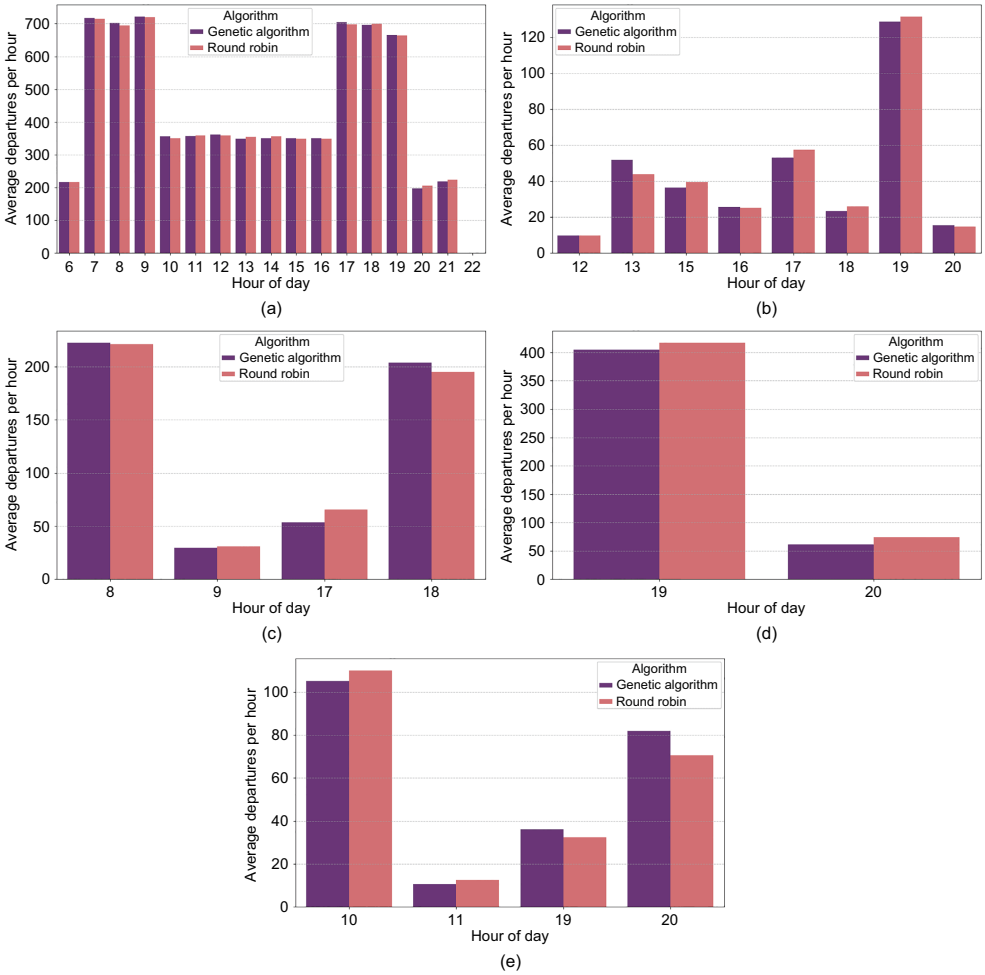


Fig. 7 Average hourly UAV throughput under five weather conditions: (a) no event; (b) moderate wind; (c) strong wind; (d) heavy rain; (e) storm. All figures share the same time axis and throughput scale

Fig. 8 provides detailed waiting time. Under normal conditions, both methods yield near-zero delays. However, in strong wind and storm scenarios, the GA sometimes fails to optimize queues, producing waiting times up to 45% longer than RR. For example, on Jan. 4 at 18:00, the GA reaches 2.2 min, compared to 1.5 min of RR.

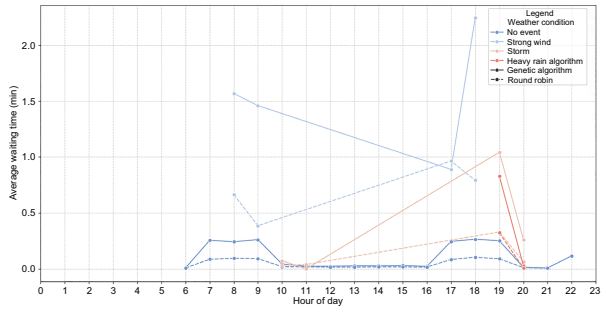


Fig. 8 Hourly average waiting time under different weather conditions

5.4 Statistical analysis and discussion

To validate the statistical properties of the output data, a normality check using the Shapiro–Wilk test and a homogeneity of variance test using Levene's test are performed, following procedures outlined by Pagotto et al. (2021). The results indicate that the data are non-parametric, as the variances in waiting times between the two methods are heterogeneous. Consequently, the Mann–Whitney U test, a robust nonparametric method suitable for large samples, is employed to assess differences between distributions (Roberge et al., 2013; Leite et al., 2021).

The Mann–Whitney U test yielded a result with a p-value of 1.76×10^{-40}, which is significantly lower than the standard threshold of 0.05. Therefore, this study rejects the null hypothesis and concludes that the waiting time distributions generated by the RR and the GA algorithms are statistically distinct. A detailed analysis of the factors contributing to the GA's underperformance is provided in Section 1 of the supplementary materials.

6 LLM-assisted GA refinement

This section presents how LLMs enhance the development of the GA through iterative reasoning, fitness redesign, and performance evaluation. The LLM (Gemini 2.5 Pro) acts as a code generator and a scientific partner that accelerates insight generation and algorithmic innovation.

6.1 Motivation and framework overview

Initial results show that RR, a standard online scheduling baseline, consistently outperforms the first version of the GA (GA v1) in terms of the average waiting time. This unexpected outcome prompts an investigation into the GA's design, particularly its fitness function, which revealed an overemphasis on pad scarcity penalties at the expense of waiting time minimization.

To address this, a structured AI-assisted development process is adopted, where the LLM plays in three roles:

1. Diagnosing failure modes through log interpretation and hypothesis generation;

2. Prototyping improved fitness functions in response to real-time user prompts;

3. Framing theoretical insights into the nature of online scheduling under uncertainty.

Fig. 9 summarizes the architecture of this hybrid optimization system.

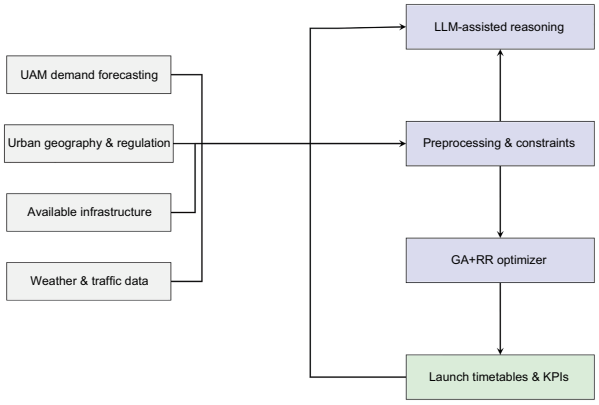


Fig. 9 Architecture of the LLM-assisted optimization framework for takeoff scheduling in vertiport

6.2 Cost function variants and prompt-based design

The LLM-assisted design yields four GA variants.

1. GA v1 (baseline and proposed by human): with original fitness function, combining the average waiting time and fixed penalties for rare pad usage.

2. GA v2 (inverted penalties): reduced weights for class 1 pads, mitigating the starvation of high-priority UAVs.

3. GA v3 (dynamic penalties): time-dependent weights $w_k(t)$ responding to traffic peaks and environmental pressure.

4. GA v4 (fairness-oriented): penalized high maximum delays, targeting worst-case user experience.

5. GA v5 (statistical optimization): minimized mean, standard deviation, and 95$^{\text{th}}$ percentile waiting times.

The design of cost function $f(c)$ evolves across these versions, as follows:

$$f(c)=\alpha_1\cdot\bar{T}_{\text{wait}}+\alpha_2\cdot\sigma_T+\alpha_3\cdot T_{95\%}+\sum_k w_k(t)\cdot T_{\text{penalty}}, \tag{6}$$

where $f(c)$ represents the total scheduling cost of a given candidate solution c and is designed to balance user experience, fairness, and system efficiency. The variables in the equation are defined as follows.

(1) $\bar{T}_{\text{wait}}$: average waiting time across all UAVs in the candidate scheduling sequence c, reflecting overall service responsiveness.

(2) σ_T: standard deviation of waiting time, capturing the dispersion or fairness among individual UAVs; lower values indicate more equitable scheduling.

(3) $T_{95\%}$: 95^{th} percentile of waiting time, representing near-worst-case delays and ensuring that no UAV experiences excessively long waits.

(4) $w_k(t)$: a time-varying penalty weight associated with launch pad k at time t, reflecting resource scarcity, priority levels, or dynamic congestion.

(5) T_{penalty}: aggregated penalty time resulting from pad contention or constraint violations (e.g., class 1 pad overuse, priority overrides, or regulatory conflicts).

(6) $\alpha_1,\alpha_2,\alpha_3$: tunable weighting coefficients for each objective term, interactively adjusted using the LLM to align the cost function with system-level performance goals, such as minimizing average delays while maintaining fairness and robustness.

Together, these components define a multi-objective optimization model, in which both user experience and infrastructure constraints are jointly considered. The LLM-assisted tuning process facilitates rapid prototyping of different weight configurations to identify Pareto-optimal trade-offs under varying demand scenarios.

6.3 Collaborative iteration with the LLM

In this study, the iterative development of the GA followed a novel human–AI collaboration methodology. Rather than using the LLM merely as a coding assistant, Gemini 2.5 Pro is employed as an interactive research partner. The methodology unfolded in three phases.

Phase 1: diagnosis and corrective analysis

Upon observing the underperformance of the GA, compared to RR, a direct dialog is initiated with the LLM. Numerical outputs and plots are supplied, prompting the LLM to analyze the performance, identify inconsistencies, and correct itself when required. For instance, "Note that you misinterpreted the dashed and solid lines. In the strong wind scenario, the dashed line (RR) remained below the solid (GA)."

The LLM acknowledges and adjusts its interpretation accordingly. This dynamic feedback loop accelerates insight generation while maintaining human oversight to ensure analytical precision.

Phase 2: hypothesis generation and rapid prototyping

Once the root issue is identified, the focus shifts to generating and testing alternative strategies. The LLM suggests new fitness function variants, incorporating fairness metrics, time-dependent weights, and statistical goals. The researcher strategically guides these discussions. For example, "Let us redesign the fitness function. Build a version that optimizes the average waiting time, standard deviation, and worst-case delay using appropriate weighting."

The LLM returns executable code and proposes novel metrics and trade-offs (e.g., 95^{th} percentile constraint). All four variants (GA v2–v5) are implemented and tested in a single working session.

Phase 3: insight extraction and theoretical framing

Finally, the LLM assists in interpreting the results. After the last test, "Review the following results: [table]. What could be going wrong again?"

The LLM translates these specifications into executable fitness logic, proposes alternatives (e.g., priority-based queuing), and justifies parameter choices. It also corrects its own misinterpretations upon clarification, demonstrating adaptable reasoning under human supervision.

All interactions are conducted through structured prompt sessions using conversational or programmatic interfaces. No fine-tuning is performed; instead, model-native few-shot and chain-of-thought prompting techniques are employed. The models are not connected to live UAV systems; their role is confined to code generation, parameter reasoning, and result interpretation within simulation-only settings.

This multi-model approach enables triangulation across different LLM reasoning styles, enhances robustness against model-specific biases, and demonstrates the feasibility of LLM-assisted

scientific workflows in engineering applications.

7 Comparative evaluation of scheduling algorithms with LLM assistance

To enhance GA performance, Gemini 2.5 Pro is employed as an active collaborator throughout an iterative refinement (GA v2–v5), followed by independent evaluation using GPT-4o Mini to generate and assess a new variant (GA v6). This approach aims to investigate both the impact of LLM guidance on performance and the robustness of AI-generated scheduling strategies across models. Detailed specifications and configuration of the LLMs are provided in Section 2 of the supplementary materials.

7.1 LLM-guided development of GA v2–v5

Table 9 summarizes the performance of all algorithm variants evaluated in this study, including the classical RR, GA v1, and four LLM-refined GA versions (v2–v5). Comparisons are based on key scheduling metrics: average waiting time, maximum waiting time, zerowait (percentage of takeoffs with no waiting), and longwait (percentage of takeoffs delayed >120 s).

A clear pattern emerges: RR consistently achieves the best worst-case performance, capping maximum waiting time at 206 s. In contrast, the baseline GA v1 exhibits poor performance under this metric, with a maximum waiting time of 738 s. These results underscore the challenges of designing effective heuristics and fitness functions without LLM-guided refinement.

Fig. 10 presents a detailed breakdown of average waiting times by UAV class (classes 1–3) and for high-priority official flights, comparing RR with GA v1–v5. While RR maintains the best overall average performance, GA v5 substantially reduces worst-case delays to 298 s, a 59.62% improvement over GA v1, and approaching the RR benchmark.

Interestingly, GA v2 and GA v3 perform slightly worse than GA v1, indicating that simple weight adjustments or dynamic penalties alone are insufficient. In contrast, GA v4 and GA v5, which incorporate fairness considerations and statistical optimization respectively, demonstrate significant improvements.

These results underscore that the alignment of the fitness function with system-level objectives, rather than algorithm complexity alone, is critical to performance. They further highlight the value of LLM guidance in navigating multidimensional design spaces and proposing effective optimization strategies.

7.2 GPT-4o: GA v6 design and evaluation

To further evaluate the robustness of LLM-assisted optimization, we conduct an independent experiment using a lightweight language model, GPT-4o Mini, to design a new fitness function for GA. This allows us to assess the quality of its output and the extent to which LLMs can independently infer domain-aligned heuristics from a one-shot prompt.

7.2.1 Generation of GA v6 using another LLM

A high-level prompt describing the scheduling problem, system objectives, and previous GA limitations is submitted to GPT-4o Mini. The model proposes a new function, calcular_fitness_v6_chatgpt, designed to reduce the average waiting time by introducing the following.

1. Dynamic penalties: Penalty weights vary hourly. During peak hours (07:00–09:00 and 17:00–19:00), delays for lower-priority flights (class 3) incur higher penalties, while delays for high-priority flights (class 1) are tolerated more to ensure the throughput of critical traffic.

2. Anti-starvation protection: Any official (class 1) flight delayed more than 180 s receives a large fixed penalty of 5000, acting as a safeguard to prevent the negligence of priority operations.

7.2.2 Numerical results and performance

The fitness function generated for GA v6 is manually integrated into the simulation environment and evaluated under the same conditions as previous GA versions. Table 10 presents a comparative analysis of GA v6 and RR, focusing on overall waiting time and quality of service (QoS) metrics.

Despite incorporating an intuitive penalty mechanism, GA v6 produces a maximum waiting time of 636 s—substantially higher than RR's 206 s, and also worse than GA v4 and GA v5 (see Eq. (6)). Although its design emphasizes prioritization and delay penalization, it lacks the statistical rigor and multi-objective balance achieved by earlier LLM-assisted iterations. A detailed comparison

Table 9 Final performance comparison: RR and GA variants (v1–v5)

Metric	RR	GA v1	GA v2	GA v3	GA v4	GA v5
Average waiting time (s)	3	10	10	10	9	9
Maximum waiting time (s)	206	738	1048	1091	393	298
Zerowait	77.39%	60.00%	60.00%	60.00%	59.99%	59.94%
Longwait	0.06%	0.31%	0.30%	0.30%	0.27%	0.34%

Table 10 Numerical comparison between the round-robin and GA v6 (GPT-4o Mini-generated)

Metric	Round robin	GA v6 (GPT-4o Mini)
Overall waiting time		
Average waiting time (s)	3	10
Maximum wait (worst case) (s)	206	636
95^{th} percentile wait (s)	24	50
QoS		
Zerowait	77.39%	60.24%
Longwait	0.06%	0.35%

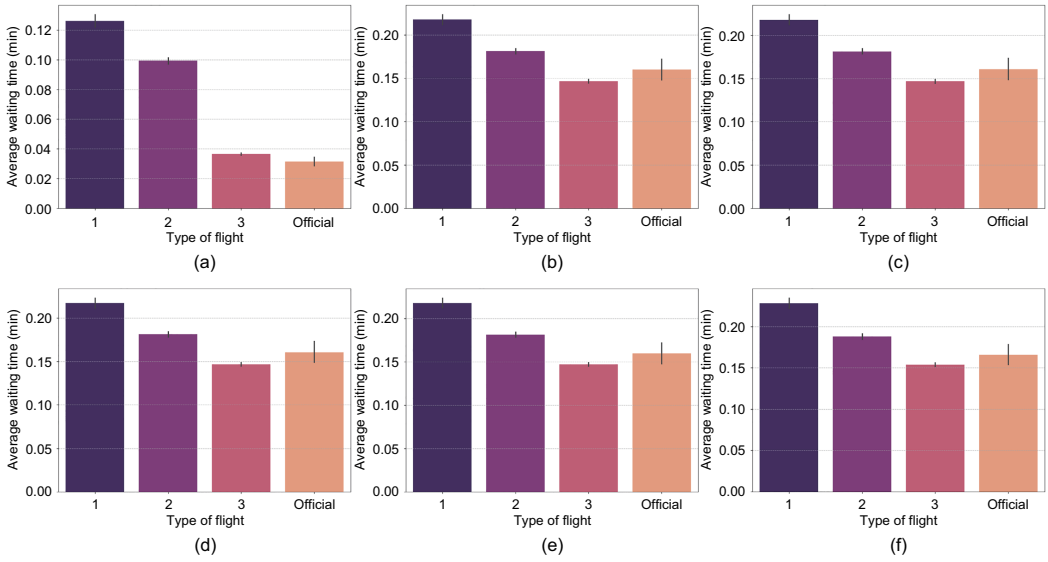


Fig. 10 Class-wise waiting time comparisons for each scheduling algorithm variant: (a) round robin; (b) GA v1 (original); (c) GA v2 (inverted); (d) GA v3 (dynamic); (e) GA v4 (fairness); (f) GA v5 (statistical)

of the fitness function objectives between GA v5 (Gemini 2.5 Pro–assisted) and GA v6 (GPT-4o Mini–assisted) is presented in Section 3 of the supplementary materials.

7.2.3 Reflections on the LLM as an independent designer

The GA v6 case demonstrates the strengths and limits of one-shot prompting with lightweight LLMs.

1. Understanding vs. Implementation: GPT-4o Mini quickly grasps the optimization goal but requires human debugging to effectively integrate its proposal.

2. Instruction vs. Iteration: Unlike the iterative development of GA v5, GA v6 is produced in a single prompt, limiting its adaptability and transparency.

This highlights the continued need for iterative, human-in-the-loop workflows to unlock the full potential of LLMs in optimization tasks.

7.3 Discussion and future directions

The results highlight several broader implications for UAM scheduling research. First, the contrast between RR and GA emphasizes the

inherent tension between reactive fairness-based schemes and model-driven optimization under static assumptions. While RR appears stronger in simplified environments, its lack of flexibility suggests limited scalability once regulatory constraints and heterogeneous UAV classes are fully integrated.

Second, the results demonstrate how LLMs can serve as both code generators and conceptual collaborators. Dialog-driven interaction with Gemini 2.5 Pro enables iterative refinements and deeper reasoning about fairness metrics, whereas GPT-4o Mini illustrates that even lightweight models can independently propose workable heuristics. These complementary interaction modes suggest that future workflows may benefit from hybrid LLM participation, combining rapid prototyping with design-level insights.

Finally, the study highlights methodological challenges in aligning fitness functions with operational realities. The difficulty of GA variants in matching RR's short-term performance underscores the need for hybrid online–offline architectures and more adaptive formulations of fairness and delay penalties. LLMs could play a key role in guiding such parameter tuning and supporting continuous algorithmic evolution.

Future research should explore (1) hybrid scheduling architectures incorporating rolling-horizon updates, (2) LLM-assisted calibration of penalty weights and thresholds, and (3) integrating these methods into real-time pipelines using live sensor and traffic data. Collectively, these directions support the development of a more adaptive, resilient, and co-evolutionary framework for UAM scheduling.

8 Conclusions

This study introduces Eixão-UAM, an LLM-assisted framework for the iterative design, simulation, and initial optimization of a structured low-altitude UAM corridor in Brasilia, focused on the Eixão axis. The framework integrates urban morphology, airspace stratification, vertiport layout, and scheduling strategies to demonstrate how architectural and regulatory constraints can be systematically translated into an operationally feasible mobility system.

The simulation results offer three key insights. First, simple online methods such as RR remain highly competitive under dynamic and uncertain conditions. Second, the effectiveness of GAs critically depends on alignment between fitness functions and system-level objectives. For example, GA v5 reduces maximum waiting time by 59.62% and average waiting time by 10% compared to GA v1, approaching the robustness of RR. Third, engaging LLMs as reasoning partners, substantially accelerates the iterative refinement of heuristics, facilitating hypothesis testing, performance diagnosis, and rapid redesign beyond conventional optimization workflows.

At the institutional level, this work is grounded in Brazil's BR-UTM ecosystem, which provides an integrated regulatory and digital framework for UAM deployment. The proposed Eixão-UAM corridor demonstrates how national governance structures can be linked with architectural and operational design principles, offering a replicable model for Latin America and beyond.

The study provides both a methodological blueprint and performance-tested case study by prototyping a vertiport-based airway and embedding LLM-guided optimization into its operation. Future research will extend beyond takeoff sequencing to address arrival management, en-route deconfliction, and resilience under live data feeds, advancing toward a fully integrated low-altitude traffic management system for Brasilia and similar metropolitan regions.

Overall, these findings demonstrate how LLM-assisted iterative design can be systematically embedded into UAM system development, offering the ATM/UAM community a novel methodological framework and an empirically validated case study that bridges conceptual UAM corridor planning with operational feasibility.

Acknowledgments

The authors acknowledge the use of generative AI tools, including ChatGPT and Gemini, which were employed to assist in formulating fitness functions and to standardize and refine the English presentation of the paper.

Contributors

Li WEIGANG conceptualized the study, developed the LLM-assisted optimization framework, and drafted the paper. Juliano Adorno MAIA implemented the RR and GA algorithms and conducted the simulations described in Section 5.

Emilia STENZEL designed the architectural layout of the Eixão-UAM corridor, and authored the corresponding text in Section 3. Lucas Ramson SIEFERT performed the LLM-based experiments presented in Section 6. All authors contributed to the discussion, revision, and finalization of the paper.

Conflict of interest

Li WEIGANG is an editorial board member of *Frontiers of Information Technology & Electronic Engineering*; he is not involved with the peer review process of this paper. All the authors declare that they have no conflict of interest.

Data availability

The data that support the findings of this study are available from the corresponding author upon reasonable request.

References

Abdellaoui R, Naser F, Velieva A, et al., 2025. Applying a comparative performance assessment framework to different airspace management concepts for urban air mobility. *CEAS Aeronaut J*, 16(3):827-847. https://doi.org/10.1007/s13272-025-00835-0

Alolaiwy M, Hawsawi T, Zohdy M, et al., 2023. Multi-objective routing optimization in electric and flying vehicles: a genetic algorithm perspective. *Appl Sci*, 13(18):10427. https://doi.org/10.3390/app131810427

ANAC, 2017a. Classes de Drones (in Portuguese). https://www.gov.br/anac/pt-br/assuntos/drones/classes-de-drones [Accessed on July 30, 2025].

ANAC, 2017b. Requisitos Gerais para aeronaves não tripuladas de uso civil (RBAC-E 94) (in Portuguese). https://www.anac.gov.br/assuntos/legislacao/legislacao-1/rbha-e-rbac/rbac/rbac-e-94 [Accessed on July 30, 2025].

Bayer DM, Bayer F, 2015. Previsão da umidade relativa do ar de brasília por meio do modelo beta autorregressivo de médias móveis. *Rev Bras Meteorol*, 30(3):319-326 (in Portuguese). https://doi.org/10.1590/0102-778620130645

Chan YY, Ng KKH, Lee CKM, et al., 2023. Wind dynamic and energy-efficiency path planning for unmanned aerial vehicles in the lower-level airspace and urban air mobility context. *Sustain Energy Technol Assess*, 57:103202. https://doi.org/10.1016/j.seta.2023.103202

Cheng AW, Witzberger KE, Isaacson DR, et al., 2022. National Campaign (NC)-1 Strategic Conflict Management Simulation (X4) Final Report. NASA/TM-2022-0018159, NASA, Moffett Field.

Cohen AP, Shaheen SA, Farrar EM, 2021. Urban air mobility: history, ecosystem, market potential, and challenges. *IEEE Trans Intell Transp Syst*, 22(9):6074-6087. https://doi.org/10.1109/TITS.2021.3082767

da Silva R, do Sul Milholi da Silva R, de Almeida Regis J, et al., 2020. Acesso ao espaço aéreo brasileiro por aeronaves não tripuladas. *Rev CIAAR*, 1(1):23-40 (in Portuguese).

DECEA, 2023. Remotely piloted aircraft systems and access to Brazilian airspace. Brazilian Airspace Control Department, ICA 100-40.

Deniz S, Wang ZB, 2024. Autonomous conflict resolution in urban air mobility: a deep multi-agent reinforcement learning approach. Proc AIAA Aviation Forum and Ascend, Article 4005. https://doi.org/10.2514/6.2024-4005

de Vasconcellos E, Regis JM, 2022. Brazil's DECEA Launches UTM Sandbox and Publishes Implementation Programme. Unmanned Airspace. https://www.unmannedairspace.info/uncategorized/brazils-decea-launches-brazils-utm-sandbox-and-published-implementation-programme [Accessed on July 30, 2025].

EASA, 2017. Drones and air mobility basics explained. https://www.easa.europa.eu/en/domains/drones-air-mobility/drones-air-mobility-landscape/basics-explained [Accessed on July 30, 2025].

FAA, 2023a. Advanced Air Mobility (AAM) Implementation Plan. https://www.faa.gov/sites/faa.gov/files/AAM-I28-Implementation-Plan.pdf [Accessed on July 30, 2025].

FAA, 2023b. Urban Air Mobility (UAM) Concept of Operations. https://www.faa.gov/sites/faa.gov/files/Urban-Air-Mobility-Concept-of-Operations-2.0.pdf [Accessed on July 30, 2025].

Ferreira DM, Rosa LP, Ribeiro VF, et al., 2014. Genetic algorithms and game theory for airport departure decision making: GeDMAN and CoDMAN. Proc 9th Int Conf on Knowledge Management in Organizations, p.3-14. https://doi.org/10.1007/978-3-319-08618-7_1

Garcia CP, Weigang L, Hirata NST, et al., 2023. ISUAM: intelligent and safe UAM with deep reinforcement learning. Proc 29th Int Conf on Parallel and Distributed Systems, p.378-383. https://doi.org/10.1109/ICPADS60453.2023.00064

Gong YW, Fan JC, Zhang RC, et al., 2025. Safe and economical UAV trajectory planning in low-altitude airspace: a hybrid DRL-LLM approach with compliance awareness. https://doi.org/10.48550/arXiv.2506.08532

Goodrich KH, Theodore CR, 2021. Description of the NASA urban air mobility maturity level (UML) scale. Proc AIAA Scitech Forum, Article 1627. https://doi.org/10.2514/6.2021-1627

Halder S, Ghosal A, Conti M, 2023. Dynamic super round-based distributed task scheduling for UAV networks. *IEEE Trans Wirel Commun*, 22(2):1014-1028. https://doi.org/10.1109/TWC.2022.3200366

ICAO, 2015. Manual on Remotely Piloted Aircraft System (RPAS). International Civil Aviation Organization, Montrèal.

Jasper FNH, Nunes AF, 2022. Soberania e controle do espaço aéreo: uma visão brasileira. *Rev Tempo Mundo*, (28):345-366 (in Portuguese). https://doi.org/10.38116/rtm28art12

Jiang YH, Li XY, Zhu GX, et al., 2023. 6G non-terrestrial networks enabled low-altitude economy: opportunities and challenges. https://doi.org/10.48550/arXiv.2311.09047

Korringa M, Snyder P, Ullrich M, et al., 2025. Optimal air corridor design for efficient integration of AAM vehicles into the NAS. Proc Integrated Communications, Navigation and Surveillance Conf, p.1-7. https://doi.org/10.1109/ICNS65417.2025.10976768

Lascara B, Spencer T, DeGarmo M, et al., 2018. Urban Air Mobility Landscape Report. McLean, VA, USA. https://www.mitre.org/publications/technical-papers/urban-air-mobility-landscape-report

Lavezzi G, Guye K, Cichella V, et al., 2023. Comparative analysis of nonlinear programming solvers: performance evaluation, benchmarking, and multi-UAV optimal path planning. *Drones*, 7(8):487. https://doi.org/10.3390/drones7080487

Leite GMC, Marcelino CG, Wanner EF, et al., 2021. Pattern classification applying neighbourhood component analysis and swarm evolutionary algorithms: a coupled methodology. Proc IEEE Congress on Evolutionary Computation, p.319-326. https://doi.org/10.1109/CEC45853.2021.9504702

Liang M, Li WG, Delahaye D, et al., 2019. Policy optimization in automated point merge trajectory planning: an artificial intelligence-based approach. Proc 38th Digital Avionics Systems Conf, p.1-8. https://doi.org/10.1109/DASC43569.2019.9081789

Liao XH, Xu CC, Ye HP, 2024. Benefits and challenges of constructing low-altitude air route network infrastructure for developing low-altitude economy. *Bull Chin Acad Sci*, 39(11):1966-1981 (in Chinese). https://doi.org/10.16418/j.issn.1000-3045.20240614002

Liu CL, Layland JW, 1973. Scheduling algorithms for multiprogramming in a hard-real-time environment. *J ACM*, 20(1):46-61. https://doi.org/10.1145/321738.321743

Liu S, Liu MM, 2025. Research on the security risk governance roadmap in low-altitude economic field based on the economic externality theory. *Eng Proc*, 80(1):14. https://doi.org/10.3390/engproc2024080014

Maia JA, Weigang L, Stenzel E, et al., 2025. Model of a corridor for urban mobility of unmanned aircraft in the "Eixão" of Brasília. Proc Brazilian Air Transport Symp (in Portuguese).

Moon H, Park J, Kim J, 2024. Development of decision support systems for regional air mobility (RAM) operations in South Korea: dynamic corridor network generation. Proc AIAA Aviation Forum and Ascend, Article 4255. https://doi.org/10.2514/6.2024-4255

Moraga Á, de Curtò J, de Zarzà I, et al., 2025. AI-driven UAV and IoT traffic optimization: large language models for congestion and emission reduction in smart cities. *Drones*, 9(4):248. https://doi.org/10.3390/drones9040248

Muna SI, Mukherjee S, Namuduri K, et al., 2021. Air corridors: concept, design, simulation, and rules of engagement. *Sensors*, 21(22):7536. https://doi.org/10.3390/s21227536

Neufert E, 2013. Neufert: Arte de Proyectar en Arquitectura. Gustavo Gili, Barcelona (in French).

Pagotto LG, Rodrigues J, Henrique FH, et al., 2021. Analysis of variance and means tests: a study applied in experiments with cotton varieties and citrumelo selections. *Brazilian Appl Sci Revi*, 5(3):1287-1296. https://doi.org/10.34115/basrv5n3-001

Pruekprasert S, Nakadai S, 2024. Enhancing safety in UAM corridors: a self-separation scheme utilizing estimated arrival times at constraint waypoints. Proc Int Workshop on ATM/CNS, p.187.

Roberge V, Tarbouchi M, Labonte G, 2013. Comparison of parallel genetic algorithm and particle swarm optimization for real-time UAV path planning. *IEEE Trans Ind Inform*, 9(1):132-141. https://doi.org/10.1109/TII.2022.2198665

Sadik AR, Ashfaq M, Mäkitalo N, et al., 2025. Urban air mobility as a system of systems: an LLM-enhanced holonic approach. Proc 20th Annual System of Systems Engineering Conf, p.1-7. https://doi.org/10.1109/SoSE66311.2025.11083807

Turchetti JV, Murça MCR, 2024. Analysis and prediction of airspace availability for urban air mobility operations in the Sao Paulo Metropolitan Region. *Transportes*, 32(1):1-20. https://doi.org/10.58922/transportes.v32i1.2896

van Nguyen T, 2020. Dynamic delegated corridors and 4D required navigation performance for urban air mobility (UAM) airspace integration. *J Aviat Aerosp Educ Res*, 29(2):57-72. https://doi.org/10.15394/jaaer.2020.1828

Varun Kumar KA, Priyadarshini R, Kathik PC, et al., 2023. Self-co-ordination algorithm (SCA) for multi-UAV systems using fair scheduling queue. *Sens Rev*, 43(4):233-242. https://doi.org/10.1108/SR-01-2022-0003

Verma S, Dulchinos V, Dan Wood R, et al., 2022. Design and analysis of corridors for UAM operations. Proc IEEE/AIAA 41st Digital Avionics Systems Conf, p.1-10. https://doi.org/10.1109/DASC55683.2022.9925820

Wang BH, Wang DB, Ali ZA, et al., 2019. An overview of various kinds of wind effects on unmanned aerial vehicle. *Meas Contr*, 52(7-8):731-739. https://doi.org/10.1177/0020294019847688

Xu L, Cao XB, Du WB, et al., 2025. Robust path planning for multiple UAVs considering position uncertainty. *Chin J Electron*, 34(4):1120-1135. https://doi.org/10.23919/cje.2024.00.015

Yang JL, Yang JF, Liu DX, et al., 2024. Competition pattern and coping strategies in near space. *Strat Study CAE*, 26(5):137-145 (in Chinese). https://doi.org/10.15302/J-SSCAE-2024.05.014

Zou LY, Munir S, Hassan SS, et al., 2024. Imbalance cost-aware energy scheduling for prosumers towards UAM charging: a matching and multi-agent DRL approach. *IEEE Trans Veh Technol*, 73(3):3404-3420.

List of supplementary materials

1 The factors contributing to the GA's underperformance

2 LLM specifications and configuration

3 Comparison of fitness function objectives between GA v5 (Gemini 2.5 Pro-assisted) and GA v6 (GPT-4o Mini-assisted)

Table S1 Comparison of fitness function objectives between GA v5 (Gemini 2.5 Pro-assisted) and GA v6 (GPT-4o Mini-assisted)

Fig. S1 Visual performance analysis of GA v6 compared to round robin