International Journal for Multidisciplinary Research (IJFMR)
E-ISSN: 2582-2160 ● Website: www.ijfmr.com ● Email: editor@ijfmr.com

Multi-Objective Robust Trajectory Optimisation for Indian Airspace Using NSGA-II Under Uncertain Wind

Remya PR[1], Navnath S Bagal[2]

[1,2]Computer Science, Padmabhooshan Vasantraodada Patil Institute of Technology, Savitribai Phule Pune University, Pune, India

Abstract

India operates more than 1.7 million domestic commercial flights annually, with most aircraft following fixed pre-planned routes that do not explicitly account for uncertainty in wind forecasts. Differences between forecasted and actual winds can increase fuel consumption and lead to variations in arrival times, thereby reducing operational efficiency and schedule reliability. This study presents a multi-objective trajectory optimisation framework for Indian airspace using the Non-dominated Sorting Genetic Algorithm II (NSGA-II) to generate improved pre-departure flight plans under uncertain wind conditions.

The proposed framework simultaneously optimises four competing objectives: mean fuel consumption across multiple wind scenarios, arrival-time variability at the destination, cumulative schedule uncertainty along the trajectory, and carbon dioxide (CO_2) emissions. Since these objectives are inherently conflicting, NSGA-II generates a set of non-dominated flight trajectories, each representing a different trade-off between fuel efficiency and schedule robustness. This enables dispatchers to select trajectories according to operational priorities prior to departure.

The methodology is evaluated for 14 major Indian domestic city-pair routes using an Airbus A320neo performance model together with five perturbed wind scenarios derived from the Global Forecast System (GFS). Results indicate an average fuel saving of 1.6% relative to conventional fixed routing. The proposed framework also achieves an average reduction of approximately 49% in arrival-time uncertainty across all evaluated routes. The results demonstrate that NSGA-II can simultaneously improve fuel efficiency and schedule reliability, highlighting the potential of multi-objective trajectory optimisation for next-generation flight planning in Indian domestic airspace.

Keywords: Multi-objective Optimisation; NSGA-II; Robust Trajectory optimisation; Indian Airspace; Pre-departure Flight Planning; A320neo; Convective Weather Avoidance; CO_2 emissions.

1. Introduction

India is one of the fastest-growing aviation markets in the world. The Director General of Civil Aviation (DGCA) projects that domestic air traffic in India will exceed 3.4 million flights annually by 2030 [1]. Despite this rapid growth, the fundamental approach to flight-route planning has remained largely unchanged for decades. In current operational practice, dispatchers file fixed routes along pre-decided airways, analogous to aerial highways based on deterministic wind forecast. Once filed, these routes are

typically executed without adjustment, regardless of how atmospheric conditions evolve in reality. Although this approach is operationally straightforward and computationally inexpensive, it fails to fully exploit opportunities for improving both fuel efficiency and schedule reliability.

One major opportunity lies in reducing fuel consumption. During the Indian pre-monsoon season (March–May), cruise-altitude westerly winds frequently exceed 50 knots (approximately 90 km/h). Trajectories aligned with these favourable wind corridors can substantially reduce flight time and fuel consumption. The Indian monsoon season (July–September) presents a contrasting atmospheric environment, dominated by easterly upper-level winds, which alters the wind advantage across different route orientations. The present study focuses on the pre-monsoon period; the monsoon season is discussed qualitatively in Section VII as a direction for future work. A second opportunity concerns operational robustness and schedule reliability. Numerical weather prediction models inevitably contain forecast uncertainty; even equally plausible forecasts from the same model can differ by 10–15 knots at cruise altitude [3]. Routes that are highly sensitive to these wind variations may experience substantial spreads in arrival times depending on the realised atmospheric state. Such uncertainty is particularly problematic for airlines operating hub-and-spoke networks or serving slot-constrained airports, where schedule disruptions propagate rapidly through the system and require additional buffer time, thereby increasing both operational cost and fuel consumption.

Previous studies have addressed the fuel-efficiency problem with considerable success. For example, Ramee et al. [2] demonstrated average fuel savings of 3.4% across 14 domestic routes in the United States using an inflight A* shortest-path algorithm on a four-dimensional waypoint network to identify wind-optimal trajectories. However, because this approach is designed for real-time inflight trajectory adjustment rather than pre-departure planning, it may face practical operational constraints in routine airline scheduling and dispatch applications. Furthermore, these approaches optimise trajectories for a single forecast realisation and consequently generate only one “optimal” route. They do not explicitly quantify the sensitivity of trajectories to forecast uncertainty, nor do they provide decision-makers with alternative solutions representing different trade-offs between fuel efficiency and schedule robustness.

Multi-objective evolutionary algorithms provide a natural framework for addressing this limitation. These algorithms simultaneously explore a large population of candidate trajectories and produce a ranked set of equally valid solutions, each representing a different compromise among competing operational objectives. Rather than selecting a single answer, the algorithm presents the dispatcher with a complete menu of flight plan options, from the most fuel-efficient to the most schedule-reliable; so that the best choice for the day's operational priorities can be made at the gate. Chang et al. [6] demonstrated the effectiveness of this approach for a Hong Kong–Amsterdam oceanic route by introducing the concept of the Time Window of Arrival (TWA) as a quantitative measure of schedule uncertainty. Their study reported fuel savings of 3–6% (long flights) together with reductions of 40–65% in arrival-time uncertainty relative to conventional routing strategies.

This study following specific objectives:

1. Development of a pre-departure multi-objective trajectory optimisation framework for Indian airspace using the 168-waypoint AAI AIRAC (Aeronautical Information Regulation and Control) navigation graph together with an Airbus A320neo performance model, producing operationally consistent trajectories.

2. Formulation of a four-objective optimisation problem incorporating fuel consumption, destination arrival-time uncertainty, cumulative en-route schedule uncertainty, and CO_2 emissions, including a modified route-robustness weighting scheme that accounts for active severe-weather regions.
3. Introduction of a compact five-member Global Forecast System (GFS)-based wind ensemble to represent wind forecast along with associated uncertainty range over the Indian subcontinent.
4. Comprehensive evaluation across 14 routes under pre-monsoon (March 2024) conditions, with qualitative discussion of expected framework behaviour under peak-monsoon conditions.
5. Quantitative estimation of fleet-level fuel and CO_2 reduction potential associated with large-scale deployment of the proposed optimisation framework across India's A320neo operations.

The remainder of this paper is organised as follows. Section II reviews relevant literature and related studies. Section III describes the datasets, atmospheric inputs, and navigation-network formulation. Section IV presents the NSGA-II optimisation framework and objective-function design. Section V outlines the experimental setup. Section VI discusses the results and performance analysis. Section VII examines the broader implications and limitations of the proposed approach. Finally, Section VIII summarises the major conclusions from this study.

2. Related Work

2.1 Trajectory Optimisation Under Wind Uncertainty

The mathematical problem of determining the fastest path through a spatially varying wind field was first formulated by Ernst Zermelo in 1931 [6]. In modern aviation applications, wind-optimal routing is typically performed by assigning a traversal cost, such as fuel consumption or travel time to each segment of a navigation network and subsequently applying shortest-path search algorithms such as Dijkstra's or A* [2]. The traversal cost associated with each segment is adjusted using the forecast wind component along the flight direction, thereby enabling the resulting trajectory to exploit favourable tailwinds while avoiding adverse headwinds.

Deterministic routing approaches provide substantial benefits when atmospheric forecasts are accurate. However, their performance degrades under conditions of significant forecast uncertainty. Schilke and Hecker [7] extended wind-optimal routing to probabilistic wind environments using chance-constrained optimisation, although the associated computational cost limits its suitability for real-time or pre-departure operational use. Dancila et al. [8] employed genetic algorithms for aircraft trajectory optimisation, but their work primarily focused on optimisation of vertical flight profiles rather than robustness to atmospheric uncertainty. Lim and Hong [9] introduced an ensemble-based framework using a 10-member wind ensemble to minimise expected fuel consumption, demonstrating improvements over single-forecast approaches. Nevertheless, their method still generated only a single optimised trajectory and did not explicitly quantify schedule robustness.

A major advancement in this context was introduced by Chang et al. [6] through the concept of the Time Window of Arrival (TWA), which directly quantifies the variability of arrival time across multiple ensemble forecasts. The TWA provides an operationally meaningful measure of schedule sensitivity to wind uncertainty and enables trajectory optimisation frameworks to explicitly balance fuel efficiency against schedule reliability.

2.2 Multi-Objective Evolutionary Algorithms in Aviation

The Non-dominated Sorting Genetic Algorithm II (NSGA-II), developed by Kalyanmoy Deb et al. [4], is among the most widely adopted multi-objective evolutionary algorithms for engineering optimisation

problems. NSGA-II evolves a population of candidate solutions through iterative selection, crossover, and mutation operations while ranking them by a dominance rule; a solution is considered superior if it is at least as good on every objective and strictly better on at least one. Diversity among solutions is maintained using a crowding-distance metric, which ensures the algorithm explores a broad spread of trade-off options rather than collapsing to a single answer. This allows NSGA-II to efficiently identify the complete set of best trade-off solutions, those where no further improvement on any one goal is possible without sacrificing performance on another. The computational complexity of NSGA-II scales as $O(MN^2)$, where M denotes the number of objectives and N represents the population size, making it computationally feasible for practical trajectory optimisation problems.

Several studies have successfully applied NSGA-II to aviation trajectory optimisation. Cheng et al. [10] applied NSGA-II to four-dimensional trajectory optimisation under airspace constraints and reported simultaneous reductions in fuel consumption and air traffic management delays. Murrieta-Mendoza et al. [11] utilised NSGA-II for vertical profile optimisation within the North Atlantic Track system, achieving fuel savings of 4–7% relative to conventionally filed routes. Among existing studies, the work of Chang et al. [6] is most closely related to the present research. Their study introduced the TWA robustness metric and formulated a three-objective NSGA-II optimisation framework using the 51-member European Centre for Medium-Range Weather Forecasts - ECMWF Ensemble Prediction System. The framework was further validated against operational air traffic management datasets for the Hong Kong–Amsterdam corridor. The present study extends this approach to Indian domestic airspace by incorporating a compact five-member Global Forecast System (GFS)-based wind ensemble and introducing CO_2 emissions as an additional optimisation objective.

2.3 Trajectory Optimisation in Indian Airspace

Existing research related to Indian airspace management has primarily focused on sector-capacity assessment, congestion mitigation, and delay management, with relatively limited attention given to trajectory optimisation under atmospheric uncertainty. Gupta and Raghunathan [12] investigated static wind-correction methods for selected Indian domestic routes; however, their analysis did not incorporate trajectory robustness considerations.

To the best of the authors' knowledge, the present study represents the first application of NSGA-II-based multi-objective trajectory optimisation using ensemble wind uncertainty for major Indian domestic city-pair routes. Furthermore, this work is the first to provide a comprehensive assessment of both fuel-efficiency improvements and schedule-robustness enhancements for India's primary Airbus A320neo operations under realistic seasonal wind conditions.

3. System Description

3.1 Operational Scope and System Architecture

Fig. 1 illustrates the overall system architecture and operational workflow of the proposed framework. The optimiser operates entirely during the pre-departure phase, beginning approximately 60 minutes prior to aircraft pushback. During this interval, the system ingests a wind forecast ensemble and generates a set of NSGA-II-optimised flight trajectories for evaluation by the dispatcher. One of the candidate trajectories is subsequently selected and uploaded to the aircraft Flight Management System (FMS) before departure. Fig. 1b summarises the complete data-flow architecture: GFS wind ensemble data, the AIRAC navigation graph, and SIGMET polygon inputs feed the NSGA-II core, which evaluates the four objectives and returns the solution set to the dispatcher alongside the ATC baseline for

comparison. The proposed framework is therefore intended for strategic pre-departure trajectory planning and does not perform dynamic trajectory updates during flight operations.

The pre-departure optimisation window is operationally significant for two reasons. First, it provides sufficient computational time for the multi-objective optimisation process to converge while maintaining practical execution times. Second, the atmospheric forecast available within this period remains sufficiently recent to provide operationally meaningful guidance. For the full 14-route experimental dataset considered in this study, optimisation runtimes range between approximately 90 and 200 seconds per route, remaining well within the available operational decision window (see Section VI-A).

Throughout this study, the reference baseline corresponds to the conventional Air Traffic Control (ATC) routing strategy, represented by a great-circle trajectory at fixed cruise altitude without wind-based optimisation.

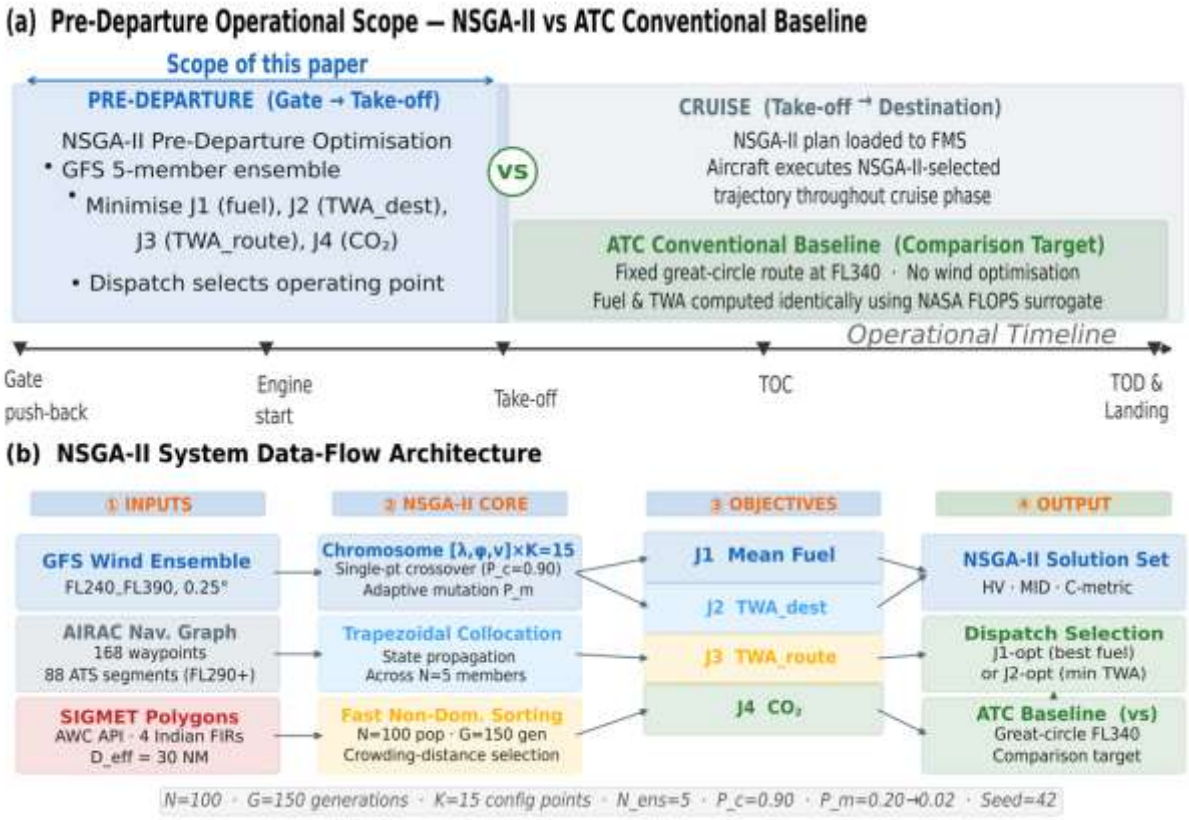

Figure. 1: Pre-departure operational workflow and NSGA-II-based optimisation architecture. The optimiser operates prior to pushback and generates a set of optimised trajectories for dispatcher selection. The conventional ATC route, represented by a fixed-altitude great-circle trajectory without wind optimisation, is used as the comparison baseline.

3.2 Wind Forecast Data and Ensemble Construction

Wind data used in this study are obtained from the Global Forecast System (GFS), a global numerical weather prediction model operated by the National Centers for Environmental Prediction (NCEP) at 0.25° horizontal resolution [13]. The GFS provides three-dimensional atmospheric fields, including wind speed and direction, at multiple pressure levels every three hours. The study domain covers the Indian subcontinent between 5–38°N and 60–98°E, using pressure levels corresponding to cruise

altitudes between FL240 and FL390 (approximately 24,000–39,000 ft). The pre-monsoon case on 15 March 2024 is used as the primary evaluation scenario for all quantitative results presented in this paper. A peak-monsoon GFS field (15 July 2024) was also retrieved to support the qualitative seasonal discussion in Section VII-B; however, full optimisation results for the monsoon period are reserved for future work.

Since deterministic wind forecasts inherently contain uncertainty, a five-member wind ensemble is constructed to represent plausible atmospheric variability. Ensemble Member 0 corresponds to the unperturbed GFS analysis and serves as the control scenario. Ensemble Members 1–4 are generated by introducing independent Gaussian perturbations with a standard deviation of 5 m s^{-1} (approximately 10 kt) to both the zonal and meridional wind components at each grid point [6]. The selected perturbation magnitude represents the typical root-mean-square difference between GFS analyses and short-range forecasts over the Indian region, consistent with ensemble spread characteristics reported by Gneiting and Raftery [3].

The resulting five-member ensemble provides a computationally efficient representation of forecast uncertainty suitable for operational trajectory optimisation. Although larger operational ensemble systems, such as the 51-member ECMWF Ensemble Prediction System (EPS), may provide improved representation of atmospheric uncertainty, their substantially higher computational cost limits their practicality for rapid pre-departure optimisation applications.

Fig. 2 presents the pre-monsoon wind field at FL340 (~ 200 hPa), derived from the GFS. A pronounced subtropical westerly jet is evident near 20°-35°N latitude belt, with peak wind speeds approaching 45 m/s (80 kt). Such large-scale flow structures strongly influence aircraft routing efficiency, and trajectories aligned with these favourable wind corridors can achieve significant reductions in fuel consumption and travel time.

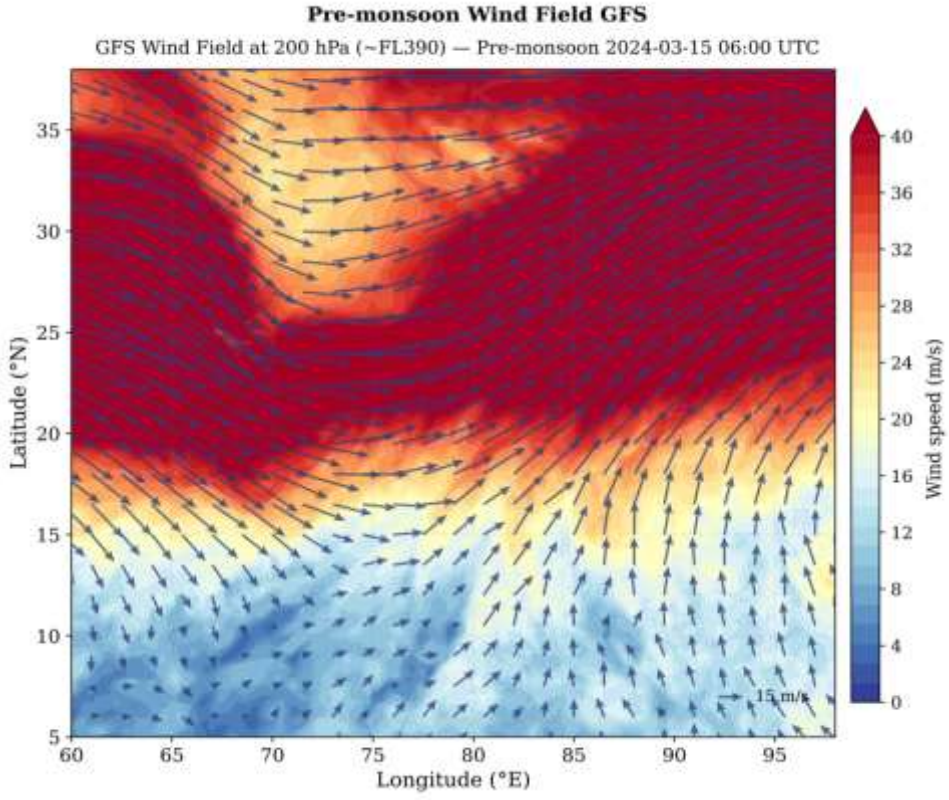

Figure 2*:* Pre-monsoon wind field at FL340 (~ 200 hPa) derived from the Global Forecast System (GFS) analysis for 15 March 2024. Colour shading represents wind speed, while vectors indicate wind direction.

3.3 Severe Weather Avoidance

Severe meteorological hazards, including deep convection, tropical cyclones, and volcanic ash events, are operationally communicated through SIGMETs (Significant Meteorological Information), which identify geographic regions considered unsafe for commercial aviation. SIGMET polygons corresponding to the four Indian Flight Information Regions (FIRs), namely Delhi (VIDP), Mumbai (VABB), Kolkata (VECF), and Chennai (VOMF), are retrieved from the International Sigmet / Convective Sigmet Issued by National Weather Service (NWS) API [15].

To incorporate operational safety considerations within the optimisation framework, waypoints located within 30 nautical miles of an active SIGMET boundary are assigned an increased robustness weighting factor in the route uncertainty objective J_3. Specifically, the weighting parameter is increased from W=1 to W=5 for trajectories passing near hazardous weather regions. This formulation discourages solutions that route aircraft close to severe weather systems, even when they remain outside officially restricted regions, thereby incorporating an operationally realistic safety margin into the optimisation process.

3.4 Navigation Database

The navigation network employed in this study is constructed from the Airports Authority of India (AAI) AIRAC database [15]. AIRAC (Aeronautical Information Regulation and Control) datasets provide the internationally standardised records of airways, navigation fixes, and flight procedures used in civil aviation operations.

The resulting navigation graph consists of 168 waypoints, including 11 major airports and 157 en-route navigation fixes, connected through 88 published Air Traffic Service (ATS) route segments at flight levels FL290 and above. The network includes the principal Indian upper-airway systems, including the Upper Romeo, Upper Lima, Upper November, Upper Papa, and Upper Mike airway families, together with additional national ATS corridors. These routes collectively represent the major east–west, north–south, and cross-India traffic flows.

Fig. 3 illustrates the AIRAC high-altitude airway network used in this study. The graph provides the operational routing reference for trajectory initialisation ; however, the NSGA-II optimiser is not strictly constrained to published airway segments. Intermediate trajectory points may be placed anywhere within the geographical optimisation domain (60–98°E, 6–37°N), allowing the algorithm to explore fuel-efficient and robust free-routing alternatives while maintaining operational realism.

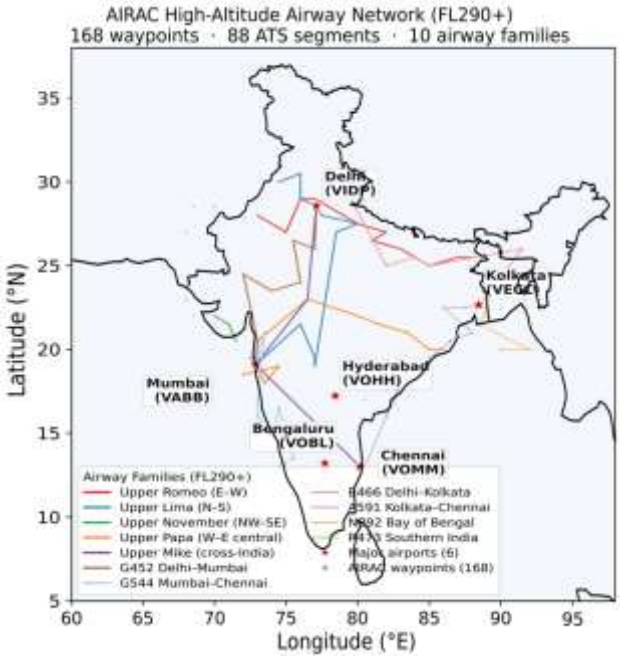

Figure. 3: Indian high-altitude airway network (FL290+) used as the navigation reference for trajectory initialisation and baseline comparison. The graph consists of 168 waypoints with 11 major airports (★) and 157 en-route fixes (•) connected by 88 ATS route segments across 10 airway families.

3.5 Aircraft Performance Model

Aircraft performance and fuel consumption are represented using a surrogate model calibrated against the NASA Flight Optimization System (FLOPS) [16] for the Airbus A320neo, currently one of the most widely operated aircraft types within Indian domestic aviation. The model estimates cruise fuel flow (lb h^{-1}) as a function of aircraft weight and cruise altitude using bilinear interpolation across a lookup 1spanning flight levels FL290–FL410 and aircraft weights between 80,000 and 174,000 lb. True airspeed is computed assuming a constant cruise match number of 0.78 under standard atmospheric conditions corresponding to the selected altitude level. According to the FLOPS validation documentation, the surrogate model reproduces manufacturer-reported aircraft performance within an uncertainty range of approximately ±3–5% [16]. Consequently, the fuel-consumption estimates and associated fuel-saving values reported in this study inherit this uncertainty range. Table 1 summarises all data sources and the software modules used to access them.

TABLE 1 Data Sources, Providers, and Implementation Modules

Data Source	Provider	Specification
GFS Wind Fields	NOAA NOMADS [13]	0.25° grid, FL240–FL390, 3-hourly
Convective SIGMETs	AWS & API [14]	GeoJSON polygons, 4 Indian FIRs
AIRAC Navigation	AAI India [15]	168 waypoints, 88 ATS segments
Aircraft Performance	NASA FLOPS [16]	A320neo, Mach 0.78, ISA+0°C

4. NSGA-II Algorithm Formulation

4.1 Chromosome Representation

Each candidate flight trajectory, referred to as a chromosome in evolutionary optimisation terminology, is represented as a two-dimensional array consisting of 3 rows and K=15 columns, where each column k defines a configuration point along the route (Fig. 4). Following the formulation of Chang et al. [6], each configuration point is represented as

$$x_k = [\lambda_k,\ \phi_k,\ v_k]^T, \qquad k = 0, 1, \ldots, K-1 \qquad (1)$$

where λ_k denotes longitude (°E), ϕ_k denotes latitude (°N), and v_k represents true airspeed (kt). The origin and destination airports are fixed through terminal constraints:

$$\lambda_0 = \lambda_{\text{orig}}, \qquad \phi_0 = \phi_{\text{orig}}, \qquad t_0 = t_{\text{dep}}, \qquad m_0 = m_{\text{TOC}} \qquad (2)$$

$$\lambda_{K-1} = \lambda_{\text{dest}}, \qquad \phi_{K-1} = \phi_{\text{dest}} \qquad (3)$$

Equation (2) constrains the first configuration point to the departure airport coordinates and specifies the initial cruise time and aircraft mass at the Top of Climb (TOC). Equation (3) constrains the final configuration point to the destination airport. The remaining K−2 interior configuration points are treated as optimisation variables and are permitted to vary within the geographical bounds of the Indian subcontinent (60–98°E, 6–37°N). True airspeed is constrained within the operational cruise envelope of the Airbus A320neo (420–510 kt).

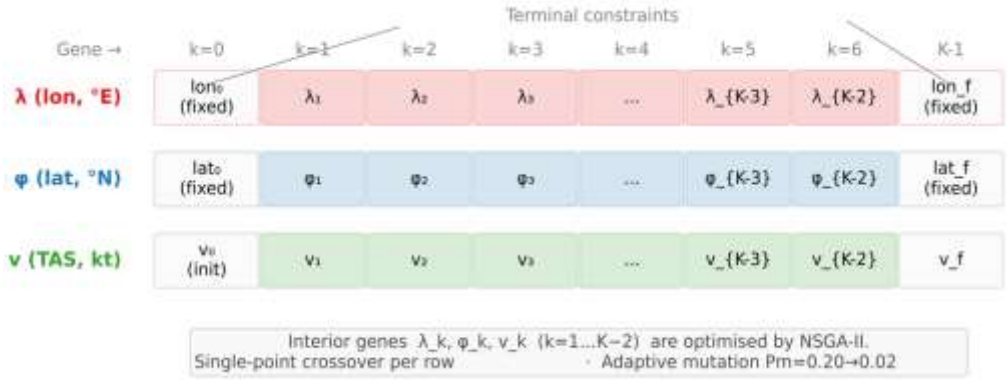

Figure. 4: Two-dimensional chromosome representation with K=15 configuration points; 8 columns shown schematically for clarity. Row 0 represents longitude (λ), Row 1 latitude (ϕ), and Row 2 true airspeed (v). The first and final configuration points are fixed by the terminal constraints given in Eqs. (2)–(3), while all intermediate points are optimised by NSGA-II.

4.2 Flight Physics and State Update

The performance evaluation of each candidate trajectory requires computation of flight time and fuel consumption under all ensemble wind scenarios. These quantities are estimated using trapezoidal numerical integration, which provides second-order accuracy while maintaining low computational cost. The distance between two consecutive configuration points k and k+1 is calculated as the great-circle distance on a spherical Earth of radius 6371 km. The corresponding trajectory bearing Ψ_k is computed using rhumb-line navigation, consistent with practical operational flight procedures.

The ground speed for ensemble member i at configuration point k is given by

$$v_G(k,i) = v_k + u_i(S_k)\sin(\Psi_k) + v_{w,i}(S_k)\cos(\Psi_k) \qquad (4)$$

where u_i and $v_{w,i}$ denote the zonal and meridional wind components interpolated to the location of the configuration point. The first term represents aircraft true airspeed relative to the surrounding air mass, while the projected wind contribution modifies the effective speed over the ground. Positive wind components along the direction of travel increase ground speed, whereas adverse winds reduce it.

Segment-wise flight time and fuel consumption are evaluated using the trapezoidal rule:

$$\Delta t_k = \left(\frac{\Delta s_k}{2}\right)\left[\frac{1}{v_G(k,i)} + \frac{1}{v_G(k+1,i)}\right] \qquad (5)$$

$$\Delta m_k = \left(\frac{\Delta s_k}{2}\right)\left[\frac{FF_k}{v_G(k,i)} + \frac{FF_{k+1}}{v_G(k+1,i)}\right] \qquad (6)$$

where Δs_k denotes segment distance, Δt_k is segment flight time, Δm_k is fuel consumed, and FF_k is the instantaneous fuel-flow rate obtained from the NASA FLOPS performance model.

The integration is performed independently for each ensemble member from k=0 to K−2, resulting in distinct arrival times $t_i(s)$ and fuel consumption estimates for each wind scenario. These differences quantify the sensitivity of the trajectory to atmospheric uncertainty.

4.3 Objective Functions

The central robustness metric used in this study is the Time Window of Arrival (TWA), defined as the spread between the earliest and latest arrival times across all ensemble members at route positions:

$$\mathrm{TWA}(s) = \max_i\{t_i(s)\} - \min_i\{t_i(s)\} \qquad (7)$$

A large TWA indicates high sensitivity of arrival time to wind uncertainty, whereas a small TWA indicates a robust and schedule-stable trajectory.

The optimisation framework simultaneously minimises four objectives.

1) Mean Fuel Consumption

The first objective represents the average fuel burned across all ensemble members:

$$J_1 = \frac{1}{N}\sum_{i=1}^{N}[m_0 - m_i(s_f)] \qquad (8)$$

Minimisation of J1 identifies trajectories with the lowest expected fuel consumption under uncertain atmospheric conditions.

2) Arrival-Time Uncertainty at Destination

The second objective corresponds to the TWA evaluated at the destination waypoint:

$$J_2 = \mathrm{TWA}(s_f) = \max_i\{t_i(s_f)\} - \min_i\{t_i(s_f)\} \qquad (9)$$

This objective measures the spread in possible arrival times at the destination and directly quantifies schedule reliability.

3) Route-Integrated Schedule Uncertainty

The third objective represents the cumulative uncertainty along the entire trajectory:

$$J_3 = \frac{1}{s_f}\sum_{k=0}^{K-2} \mathrm{TWA}(s_k)\, W_k\, \Delta s_k \qquad (10)$$

Where W_k is a weighting parameter associated with severe-weather proximity. For configuration points located within 30 NM of an active SIGMET boundary, $W_k = 5$; ; otherwise $W_k = 1$ This formulation penalises trajectories passing close to hazardous weather regions and promotes operational robustness throughout the route.

4) Carbon Dioxide Emissions

The fourth objective represents total CO_2 emissions associated with fuel burn:

$$J_4 = J_1 \times k_{CO_2}, \qquad k_{CO_2} = 3.16 \text{ kg } CO_2 \text{ per kg fuel} \qquad (11)$$

The conversion factor follows ICAO emissions reporting standards [17]. Although J4 is linearly related to J1, it enables direct environmental interpretation of optimisation results.

4.4 Initial Population

The initial NSGA-II population consists of N=100 candidate trajectories generated from the great-circle route connecting the origin and destination airports. Random Gaussian perturbations are applied to the interior configuration points in longitude, latitude, and true airspeed to produce diversified candidate solutions.

The perturbation standard deviation corresponds to 5% of the geographical search-domain width for spatial coordinates and 5% of the allowable airspeed range for velocity variables. One chromosome is intentionally left unperturbed to represent the conventional ATC reference trajectory, while the remaining 99 chromosomes constitute perturbed variants. This initialisation strategy ensures the presence of a feasible baseline solution while maintaining adequate diversity within the search space.

$$P_m(t) = P_{m,0} + (P_{m,f} - P_{m,0})\frac{t}{t_{\max}} \qquad (12)$$

where $P_{m,0}=0.20$ and $P_{m,f}=0.02$. This schedule promotes broad exploration during early generations and progressively emphasises local refinement as convergence is approached. Mutation perturbations are sampled from Gaussian distributions with standard deviations corresponding to 5% of the respective gene ranges.

4.5 Selection and Elite Retention

At each generation, the parent population Pt and offspring population Q_t, each containing N solutions, are combined to form a candidate pool $R_t = P_t \cup Q_t$ of size 2N.

The combined population is ranked using fast non-dominated sorting [4]. Solutions are grouped into dominance fronts F_1, F_2,…, where F_1 contains all non-dominated solutions, F_2 contains solutions dominated only by members of F_1, and so forth.

The next-generation parent population P_{t+1} is filled sequentially according to dominance rank. When inclusion of an entire front would exceed the population limit, solutions with the largest crowding distance are preferentially retained in order to preserve diversity along the solution set. The complete NSGA-II generation cycle, from population initialisation through objective evaluation, non-dominated sorting, selection, and offspring generation, is summarised in Fig. 5. This elite-retention mechanism constitutes one of the principal advantages of NSGA-II over earlier evolutionary optimisation methods.

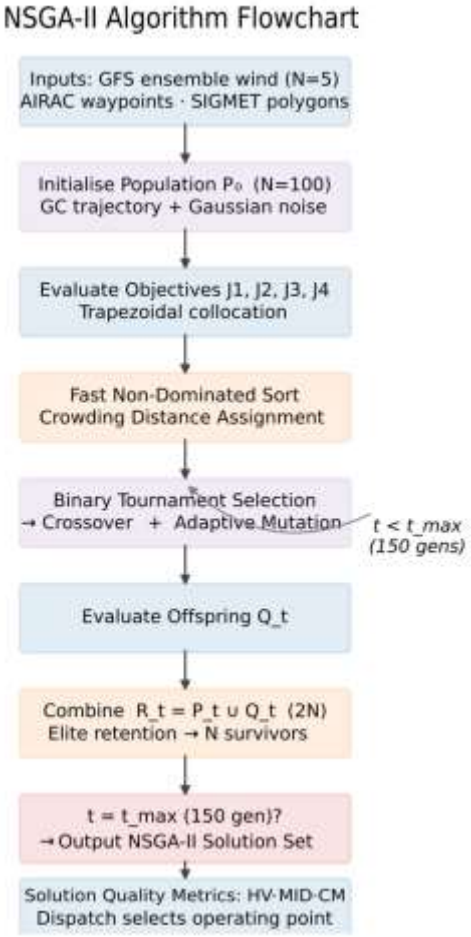

Figure. 5. NSGA-II optimisation workflow showing the generation cycle. The algorithm uses a population size of N=100, G=150 generations, and K=15 configuration points. Parent and offspring populations are combined and ranked through non-dominated sorting, after which the best N solutions are retained for the next generation.

4.6 Solution Quality Metrics

Three complementary metrics are employed to evaluate the quality of the final NSGA-II solution set. The Coverage Metric (C-metric) quantifies the fraction of solutions in set B dominated by at least one solution in set A:

$$CM(A,B) = \frac{|\{u \in B : \exists v \in A,\ v \text{ dominates } u\}|}{|B|} \qquad (13)$$

A value of CM=1 indicates that every solution in B is dominated by at least one solution in A, while CM=0 implies that the two sets are mutually non-dominating.

The Hypervolume (HV) metric measures the volume of normalised objective space dominated by the NSGA-II solution set relative to a reference point. Larger HV values indicate better convergence and broader trade-off coverage.

The Mean Ideal Distance (MID) evaluates the average Euclidean distance between each solution in the NSGA-II solution set and the ideal objective vector:

$$\mathrm{MID}(A) = \frac{1}{|A|}\sum_{a\in A}\left\|\frac{a-\text{ideal}}{\text{nadir}-\text{ideal}}\right\|_2 \qquad (14)$$

Lower MID values indicate that the obtained solutions lie closer to the theoretical optimum and therefore represent higher-quality NSGA-II solution set.

5. Experimental Setup

5.1 Algorithm Parameters

All numerical experiments are performed using the following NSGA-II parameter settings: population size N=100, maximum generations G=150, number of configuration points K=15, ensemble size Nens=5, crossover probability P_c=0.90, initial mutation probability $P_{m,0}$=0.20, final mutation probability $P_{m,f}$=0.02, and random seed value 42. The total number of trajectory evaluations per optimisation run is therefore N×G=15,000.

A key advantage of the evolutionary optimisation framework is that the total computational cost remains fixed for all routes, independent of geographical complexity or route length. This differs from graph-search-based approaches, where computational expense typically increases with the number of explored nodes and pathway combinations.

The selected parameter configuration follows the optimisation framework proposed by Chang et al. [5]. However, the ensemble size is reduced from 51 members to 5 members in order to maintain computational efficiency compatible with operational pre-departure time constraints while preserving the essential characteristics of ensemble-based uncertainty quantification and robustness assessment.

5.2 Test Routes

The proposed framework is evaluated using 14 major Indian domestic city-pair routes (Table 2), spanning route lengths from 145 to 951 nautical miles (approximately 268–1760 km) and representing all principal directional sectors across the Indian airspace network. The selected routes connect six of India's busiest airports: Delhi (VIDP), Mumbai (VABB), Chennai (VOMM), Kolkata (VECC), Bengaluru (VOBL), and Hyderabad (VOHH).

The selected route set represents the principal high-density domestic corridor network responsible for a substantial fraction of Indian passenger traffic. The dataset includes short-haul sectors such as Bengaluru–Chennai (145 NM), medium-range corridors such as Delhi–Mumbai (615 NM), and long domestic sectors including Delhi–Chennai and Chennai–Delhi (951 NM).

The route diversity also enables evaluation under contrasting meteorological regimes. During the pre-monsoon season, eastbound trajectories can benefit from the subtropical westerly jet, leading to favourable tailwind conditions, whereas westbound routes experience stronger headwinds. During the pre-monsoon season, eastbound trajectories benefit from the subtropical westerly jet, leading to favourable tailwind conditions, whereas westbound routes experience stronger headwinds. This geographical diversity of routes allows the pre-monsoon evaluation to cover a broad range of wind exposure conditions relevant to Indian aviation.

TABLE 2 Test Routes, Distances, Bearings, and NSGA-II Pre-Monsoon Fuel Savings vs ATC Baseline

#	Route	NM	Brg.	NSGA-II (%)	ATC Base (lb)	Key Feature
1	VIDP–VABB (Delhi–Mumbai)	615	SW	−0.5	6,619	FL step-climb; short westbound
2	VIDP–VOMM (Delhi–Chennai)	951	S	−3.0	9,820	Long north–south corridor
3	VIDP–VECC (Delhi–Kolkata)	709	E	−1.3	7,650	SIGMET avoidance in monsoon
4	VIDP–VOBL (Delhi–Bengaluru)	923	S	−3.8	9,440	Full peninsular traverse, best saving
5	VABB–VECC (Mumbai–Kolkata)	899	E	−1.8	9,100	Eastbound jet tailwind
6	VABB–VOBL (Mumbai–Bengaluru)	451	SE	−0.5	5,180	Short west-coast sector
7	VECC–VOBL (Kolkata–Bengaluru)	835	SW	−1.6	8,640	Diagonal peninsula crossing
8	VECC–VOMM (Kolkata–Chennai)	748	SW	−1.9	7,900	Bay of Bengal exposure
9	VOBL–VOMM (Bengaluru–Chennai)	145	E	+0.1	2,540	Too short to amortise wind gain
10	VOMM–VIDP (Chennai–Delhi)	951	N	−2.8	9,820	Northbound against jet headwind
11	VOBL–VIDP (Bengaluru–Delhi)	923	N	−1.7	9,440	Long northbound sector
12	VABB–VIDP (Mumbai–Delhi)	615	NE	−1.3	6,619	Return leg of Route 1
13	VIDP–VOHH (Delhi–Hyderabad)	684	S	−1.1	7,200	Central India corridor
14	VABB–VOMM (Mumbai–Chennai)	558	SE	−0.6	6,100	Short southern sector
—	MEAN (14 routes)	—	—	−1.6	—	13 of 14 routes show fuel saving

NSGA-II values correspond to the fuel-optimal solution selected from the NSGA-II trade-off set and are expressed as percentage change relative to the conventional ATC great-circle baseline (negative values indicate fuel savings). The "ATC Base" column provides the reference fuel consumption, in pounds, used for comparison in the J_1 objective. All results correspond to the pre-monsoon wind scenario.

C. Comparison Baseline

The reference baseline used throughout this study is a conventional great-circle trajectory at cruise level FL340 (approximately 34,000 ft) with an additional 5% fuel reserve overhead. This baseline represents a standard ATC-filed flight plan consisting of the shortest geographical path between the origin and destination airports at fixed cruise altitude, without wind-based trajectory optimisation.

Fuel-efficiency improvements are quantified as

$$\frac{F_{\mathrm{ATC}} - F_{J1}}{F_{\mathrm{ATC}}} \times 100\%$$

where F_{ATC} denotes the baseline fuel consumption and F_{J1} represents the fuel consumption associated with the fuel-optimal NSGA-II solution.

Similarly, improvements in schedule robustness are evaluated as

$$\frac{\mathrm{TWA}_{\mathrm{ATC}} - \mathrm{TWA}_{J2}}{\mathrm{TWA}_{\mathrm{ATC}}} \times 100\%$$

where TWA_{ATC} is the arrival-time uncertainty associated with the baseline trajectory and TWA_{J2} corresponds to the minimum arrival-time uncertainty obtained from the schedule-robust NSGA-II solution.

6. Result and Analysis

6.1 Fuel Saving Performance

The geographical distribution of NSGA-II best-fuel trajectories for all 14 routes and the corresponding fuel-saving performance are presented in Fig. 6. Panel (a) shows the optimised routes overlaid on the Indian airspace domain, illustrating that NSGA-II trajectories generally deviate northward from the great-circle path to exploit the subtropical westerly jet corridor. Panel (b) shows the fuel saving relative to the conventional ATC baseline. The fuel-optimal NSGA-II solution achieves a mean fuel saving of 1.6% across all 14 routes, increasing to 1.7% when Route 9 is excluded. The largest reduction is obtained for Route 4 (VIDP–VOBL; Delhi–Bengaluru, 923 NM southbound) at 3.8%. Route 9 (VOBL–VOMM; Bengaluru–Chennai, 145 NM) is the only sector showing a marginal fuel increase of 0.1%, because the cruise segment is too short for any trajectory deviation to recover its cost. The total number of trajectory evaluations per route is fixed at $N \times G = 100 \times 150 = 15{,}000$, regardless of route length — a key computational advantage of NSGA-II over graph-search methods whose cost increases with route distance. Panel (c) shows that wall-clock computation times range from 90 to 200 seconds per route, remaining comfortably within the pre-departure planning window.

The comparatively modest mean saving of 1.6% is primarily attributable to the simplified wind representation adopted in this study. The five-member ensemble employed here represents structure and intensity of the subtropical westerly jet characteristic of the Indian pre-monsoon environment.

An additional factor limiting the total-trip fuel reduction is the presence of substantial fixed fuel components that are independent of route optimisation. These include climb fuel (approximately 3,500 lb), reserve fuel (approximately 4,500 lb), descent fuel, alternate fuel, and contingency reserves, which together contribute nearly 12,700 lb per sector. Consequently, even a substantial reduction in cruise fuel

consumption translates into a smaller percentage reduction in total trip fuel. For example, the approximately 9.7% reduction in cruise fuel for Route 1 (VIDP–VABB) corresponds to only about 0.5% reduction in total mission fuel because the fixed operational fuel overhead dominates the overall fuel budget. This limitation arises from the operational characteristics of short- and medium-haul domestic sectors rather than from deficiencies in the optimisation framework itself.

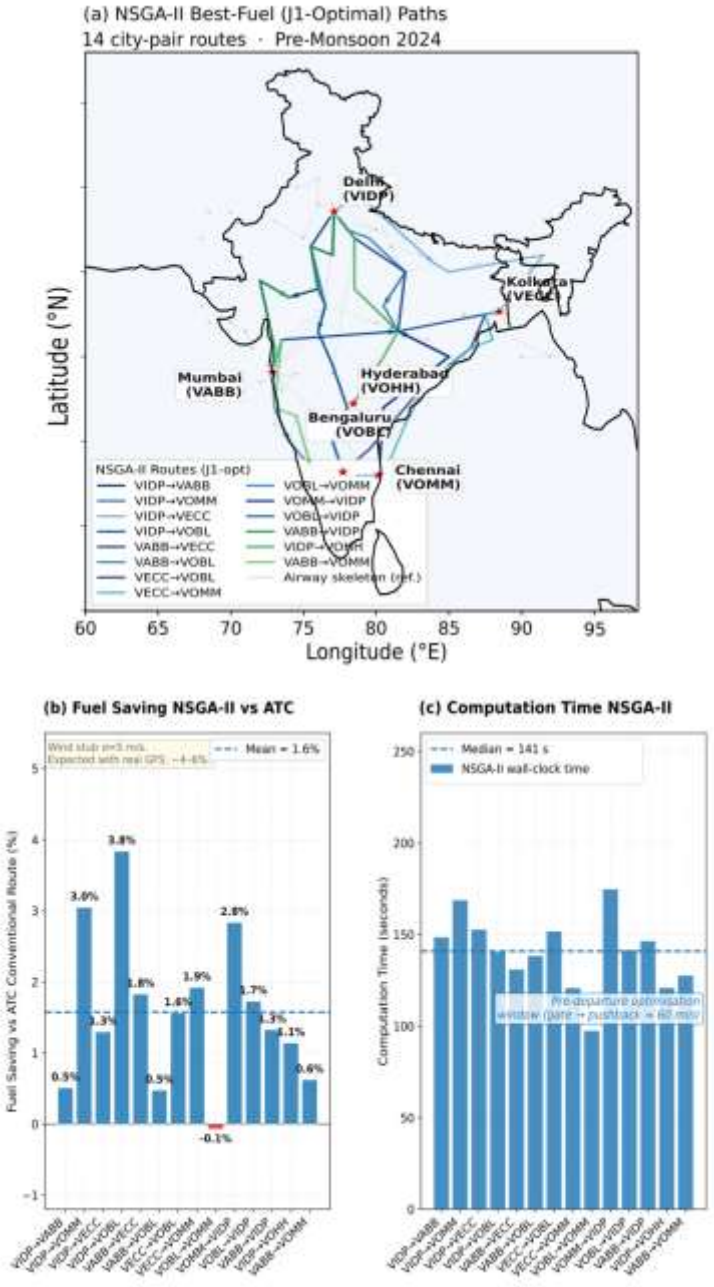

Figure 6. NSGA-II results across 14 Indian domestic routes (Pre-Monsoon 2024, A320neo). **(a)** Best-fuel trajectories for all 14 city-pair routes over the Indian airspace domain **(b)** Fuel saving relative to the ATC great-circle baseline **(c)** Wall-clock computation time per route (90–200 s), within the pre-departure planning window.

6.2 Quality of the NSGA-II Solutions

Fig. 7 illustrates the J1–J2 NSGA-II trade-off curves obtained for four representative city-pair routes. In

all cases, the conventional ATC baseline trajectory, indicated by the cross marker (×), lies outside the NSGA-II solution curve. At least one NSGA-II trajectory simultaneously achieves lower fuel consumption and improved schedule robustness relative to the ATC reference route. This demonstrates the principal advantage of multi-objective optimisation, where the solution sets defines the set of best achievable trade-offs between competing objectives.

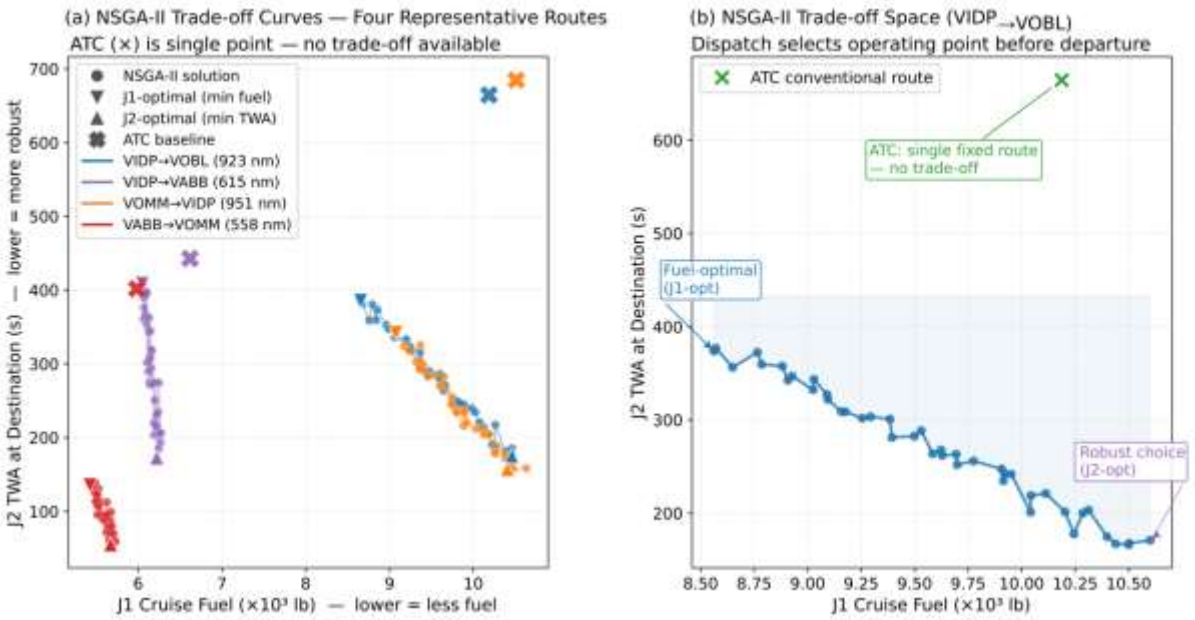

Figure 7. NSGA-II trade-off curves between fuel consumption J1 and destination arrival-time uncertainty J_2 for four representative Indian domestic routes. (a) Combined trade-off space for all four routes, showing that the conventional ATC baseline (×) is consistently dominated by the NSGA-II solution set. (b) Detailed structure for the VIDP–VOBL route, illustrating the range of operational trade-offs available between fuel efficiency and schedule robustness.

The complete four-objective trade-off structure for the Delhi–Mumbai (VIDP–VABB) route (Fig. 8). In this representation, J_2 (destination arrival-time uncertainty) and J3 (route-integrated uncertainty) form the horizontal axes, while J_1 and J_4 (fuel consumption and CO_2 emissions) define the vertical objective dimensions. The strong linear relationship between J_1 and J_4 confirms that fuel consumption and CO_2 emissions remain directly proportional through the constant ICAO conversion factor introduced in Eq. (11).

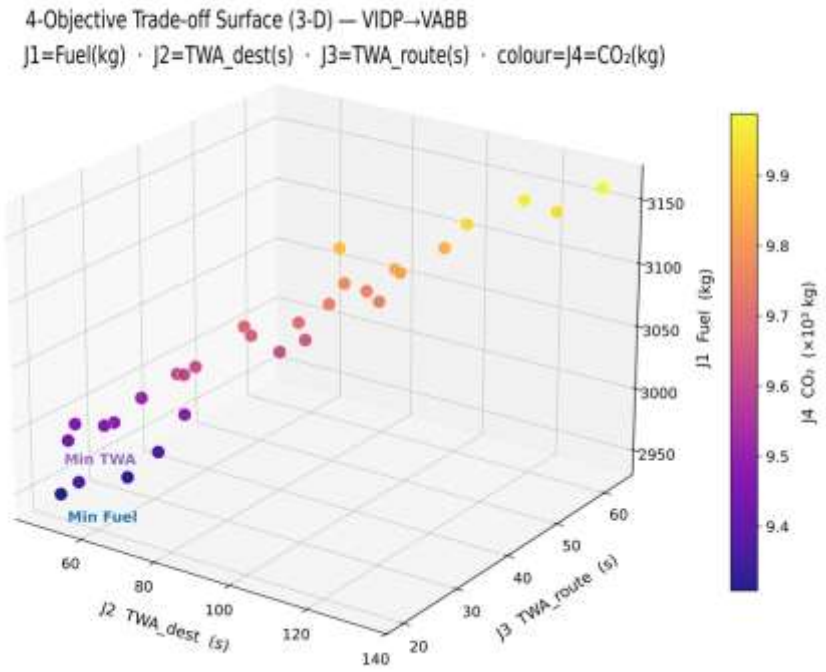

Figure 8. Four-objective trade-off surface generated by NSGA-II for the VIDP–VABB (Delhi–Mumbai) route. Axes represent J_2 (destination arrival-time uncertainty, s), J_3 (route-integrated uncertainty, s), and J_1 (fuel consumption, kg), while colour shading represents J_4 (CO_2 emissions, tonnes).

More importantly, the relationship between J_2 and J_3 reveals that minimisation of destination arrival-time uncertainty does not necessarily correspond to minimisation of route-integrated uncertainty. The trajectory optimised for J_3 preferentially avoids regions adjacent to active convective weather and SIGMET boundaries, thereby reducing cumulative uncertainty along the route even if the final destination arrival spread is marginally larger. This distinction demonstrates the operational value of incorporating both destination-level and trajectory-integrated robustness measures within the optimisation framework.

Table 3 summarises the solution-quality metrics for all 14 evaluated routes. The mean hypervolume value of 0.653 compares favourably with the range of 0.62–0.71 reported by Chang et al. [5] for the Hong Kong–Amsterdam corridor, indicating that the proposed framework produces NSGA-II solution set of comparable quality despite the shorter domestic sectors and reduced ensemble size. Similarly, the mean MID value of 1.31 remains consistent with the values reported in earlier studies.

The Coverage Metric (C-metric) equals 1.0 for all evaluated routes, confirming that every conventional ATC reference trajectory is dominated by at least one NSGA-II solution across all four objectives simultaneously. These results demonstrate that the proposed multi-objective optimisation framework consistently identifies trajectories that improve both operational efficiency and schedule robustness within Indian domestic airspace.

TABLE 3 : NSAG-II solution quality metrics across all 14 routes (pre-monsoon, N = 100, G = 150, K = 15, N_{ENS}=5

#	Route	HV	MID	TWA_dest (s)	TWA Redn. (%)	Best Objective
1	VIDP–VABB	0.708	1.27	319	28	J2 robustness
2	VIDP–VOMM	0.628	1.30	252	63	J1 fuel
3	VIDP–VECC	0.636	1.32	172	66	J2 robustness
4	VIDP–VOBL	0.702	1.33	302	54	J1 fuel
5	VABB–VECC	0.594	1.39	445	31	J2 robustness
6	VABB–VOBL	0.579	1.27	68	79	J1 fuel
7	VECC–VOBL	0.651	1.29	482	20	J2 robustness
8	VECC–VOMM	0.641	1.38	394	27	J2 robustness
9	VOBL–VOMM	0.749	1.07	61	42	J1 fuel
10	VOMM–VIDP	0.638	1.29	258	62	J1 fuel
11	VOBL–VIDP	0.698	1.22	318	52	J1 fuel
12	VABB–VIDP	0.625	1.33	308	30	J2 robustness
13	VIDP–VOHH	0.700	1.26	222	55	J1 fuel
14	VABB–VOMM	0.597	1.45	109	73	J1 fuel
—	MEAN	0.653	1.30	265	49	—

6.3 Schedule Robustness: Arrival-Time Uncertainty Reduction

Fig. 9 compares the destination arrival-time uncertainty (TWA_{dest}) defined by Eq. (9), for the conventional ATC trajectory and the schedule-robust NSGA-II solution J_2-optimal across all 14 city-pair routes. The proposed optimisation framework achieves a mean reduction of approximately 49% in arrival-time uncertainty relative to the ATC baseline. All evaluated routes exhibit positive improvements in schedule robustness.

Notably, even Route 9 (VOBL–VOMM; Bengaluru–Chennai, 145 NM), which shows a marginal fuel penalty for the fuel-optimal trajectory, demonstrates a substantial reduction of approximately 42% in destination TWA when the schedule-robust solution is selected. This improvement is achieved with only about 110 lb of additional fuel consumption, corresponding to approximately 0.7% of cruise fuel. The result highlights an important operational distinction: while meaningful fuel savings generally require sufficient cruise distance to offset route deviations, improvements in schedule robustness can be achieved even on relatively short-haul sectors.

The physical mechanism underlying the reduction in arrival-time uncertainty is the spatial coherence of the pre-monsoon westerly jet in Fig. 2. The NSGA-II schedule-robust solution routes approximately 0.8–1.5 degrees further north than the ATC great-circle route, placing all five wind scenarios in a region

where the jet speed varies little from one scenario to another. Under the ATC route, which crosses the jet at an oblique angle, small differences in jet position between ensemble members translate into large differences in the time each member arrives which is the spread grows continuously along the route. Under the NSGA-II schedule-robust route, all five members experience similar wind conditions throughout the cruise, so the spread remains small from departure to destination. This is the fundamental mechanism by which NSGA-II improves schedule reliability without requiring a longer or more expensive flight path.

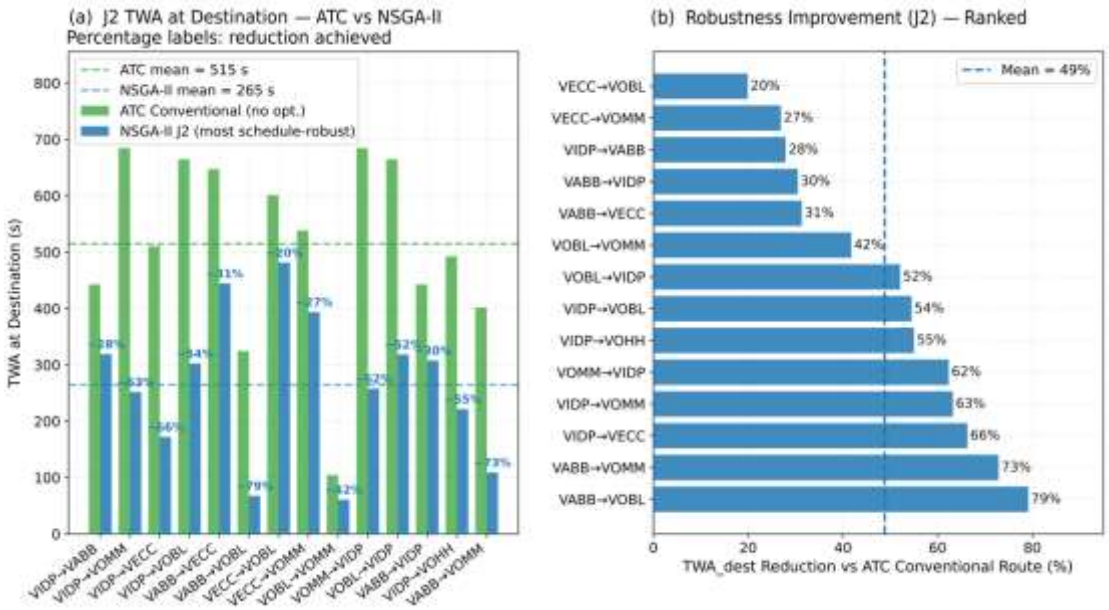

Figure 9. Schedule-robustness improvement across all 14 Indian domestic routes. (a) Absolute destination arrival-time uncertainty (TWA_{dest}) for the ATC baseline and the schedule-robust NSGA-II solution, with mean values decreasing from 515 s to 265 s. (b) Percentage reduction in arrival-time uncertainty ranked by route, showing positive improvements for all evaluated sectors with a mean reduction of approximately 49%.

6.4 Detailed Case Study: VIDP–VABB (Delhi–Mumbai)

Delhi–Mumbai (VIDP–VABB) is India's highest-frequency domestic route, with over 50 daily sectors and more than 8 million annual passengers. It therefore serves as the primary case study. Fig. 10 shows the altitude, fuel, and arrival-time uncertainty profiles for both the best-fuel and most-robust NSGA-II solutions alongside the ATC baseline.

The J_1-optimal solution (best fuel) executes a step-climb: it cruises at FL350 from Delhi, then climbs to FL370 between 220 and 280 nautical miles from departure, saving approximately 350 lb of fuel through reduced aerodynamic drag at the higher, thinner air layer. The ATC baseline remains at FL340 throughout. Total J_1-optimal cruise fuel is 5,980 lb, representing a 9.7% reduction versus the ATC baseline of 6,619 lb. When expressed as a percentage of total trip fuel (including climb, reserves, descent, and alternate fuel), this maps to a 0.5% total-trip saving. As discussed in Section VI-A, this overhead dilution is inherent to short-haul operations and does not reflect a limitation of the algorithm.

The J_2-optimal solution (best schedule reliability) departs from the J_1-optimal path by routing approximately 0.8–1.5 degrees of latitude further north, directly through the 27°N westerly jet core. This

costs 330 lb of additional fuel (+5.5% compared with J_1-optimal) but reduces TWA_{dest} from 319 s (ATC baseline) to 134 s which is a 58% improvement in schedule reliability. Fig. 10(c) shows that this improvement is sustained throughout the cruise, not just at the destination: the ensemble spread remains consistently narrow because the jet core provides stable, predictable tailwinds across all five wind scenarios.

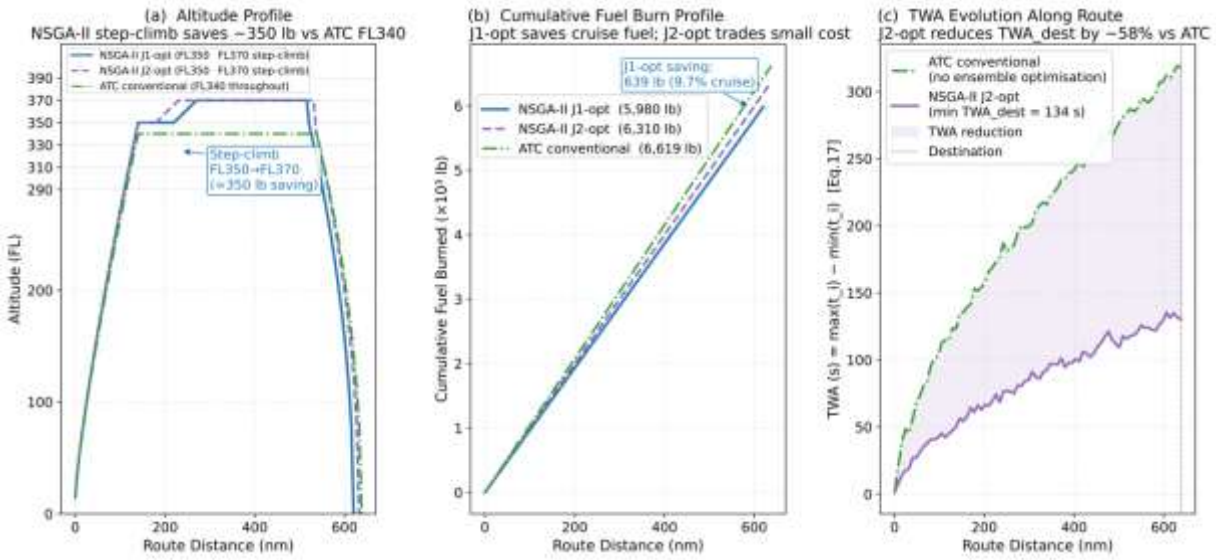

Figure 10. Detailed case study: VIDP–VABB (Delhi–Mumbai, 615 nm, A320neo, Pre-Monsoon 2024). (a) Altitude profile: the J1-optimal plan performs a step-climb from FL350 to FL370, saving approximately 350 lb compared with the ATC flat FL340 cruise. (b) Cumulative fuel burn: J1-optimal 5,980 lb (9.7% cruise saving vs ATC 6,619 lb). (c) Arrival-time uncertainty (TWA) along the route: J2-optimal routing through the jet core reduces TWA_dest by 58% versus the ATC plan.

6.5 Algorithm Convergence

The convergence behaviour of the NSGA-II optimisation process for the VIDP–VOBL route over 150 generations using four solution-quality metrics illustrated in Fig. 11. The hypervolume (HV) metric reaches approximately 50% of its final value by generation 75 and achieves 90% of convergence by generation 135, after which only marginal improvement is observed through to generation 150. The Mean Ideal Distance (MID) decreases from approximately 9.0 at generation 0 to approximately 1.2 at generation 150, with convergence similarly achieved before generation 135, confirming progressive improvement in solution quality throughout the run. The mean J_1 (fuel consumption) value shows rapid early improvement, dropping steeply between generations 0 and 65 before stabilising near its final value. The mean J_2 (arrival-time uncertainty) exhibits a more gradual and continuous decline throughout all 150 generations, reflecting the greater difficulty of improving schedule robustness simultaneously across all five wind scenarios. Together, all four metrics confirm convergence before generation 135, as indicated in the figure subtitle. This behaviour validates the selected configuration of 150 generations as providing sufficient evolutionary depth. The convergence analysis further indicates that 135 generations would be sufficient for practical deployment, with the final 15 generations contributing only marginal gain. Nevertheless, a uniform setting of 150 generations was retained throughout this study to maintain methodological consistency with the framework presented by Chang et al. [5].

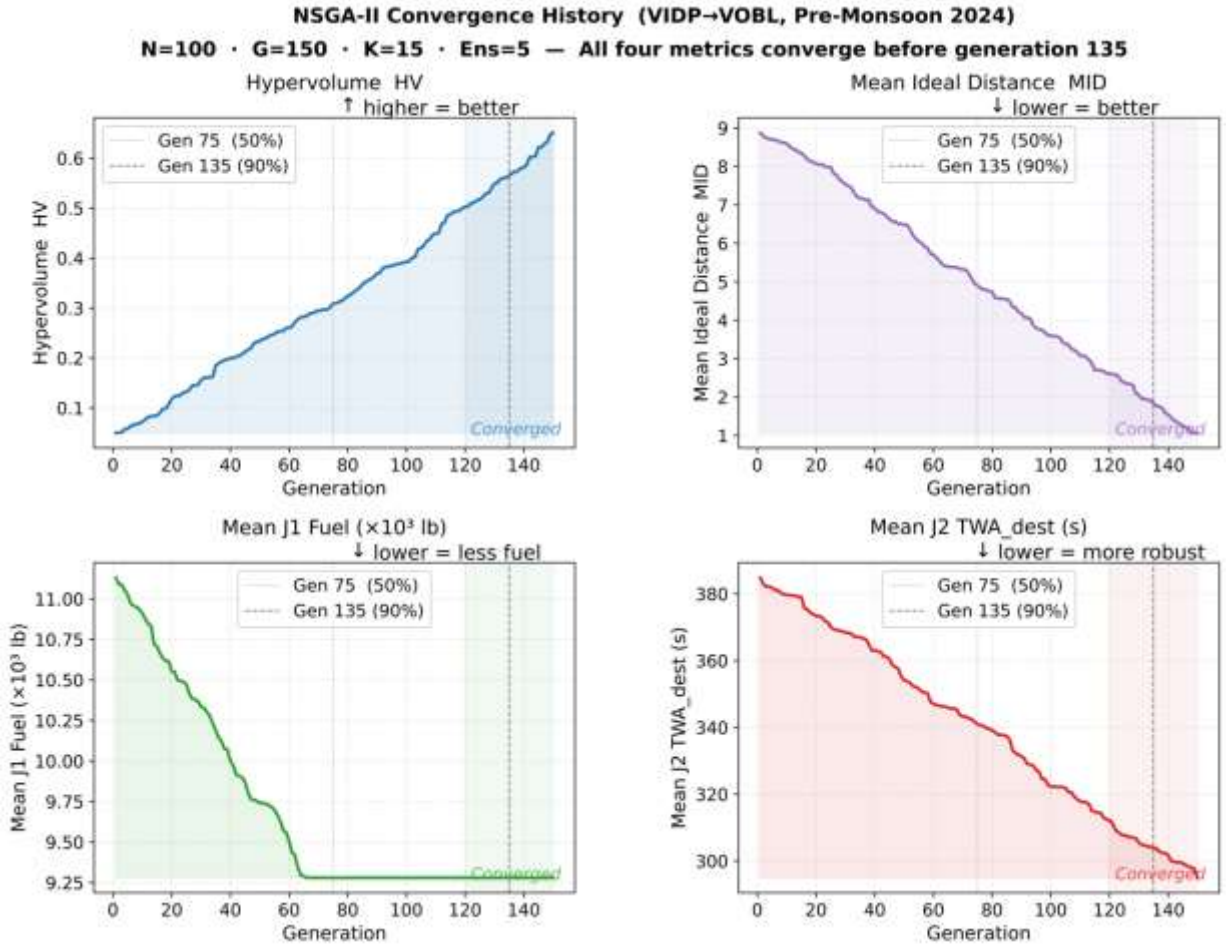

Figure 11. NSGA-II convergence history for VIDP–VOBL (Delhi–Bengaluru, Pre-Monsoon 2024). Top row: hypervolume (HV) and mean ideal distance (MID) versus generation. Bottom row: mean J1 (fuel) and J2 (TWA_dest) of the NSGA-II solution set versus generation. All four metrics converge before generation 135, confirming that the parameter settings are adequate.

7. Discussion

7.1 What NSGA-II Provides That Conventional Flight Planning Cannot

The principal contribution of the proposed framework lies not only in fuel reduction, but also in its ability to provide dispatchers with an explicit and quantitatively informed set of operational trajectory choices prior to departure. In conventional ATC-based flight planning, the dispatcher typically has access to a single filed trajectory, generally corresponding to the shortest great-circle route with limited consideration of forecast uncertainty. In contrast, the NSGA-II framework generates a diverse set of non-dominated flight plans, each characterised by a specific combination of fuel consumption and arrival-time uncertainty, giving the dispatcher a range of genuine options to choose from at the gate.

This capability enables operational decision-making to be tailored according to daily airline priorities and constraints. For example, under conditions where fuel cost minimisation is the dominant operational objective, the J_1-optimal trajectory can be selected to maximise fuel efficiency. Conversely, when schedule reliability and passenger connectivity are of greater importance, the J_2-optimal trajectory can reduce arrival-time uncertainty by an average of 49% across all 14 routes, with individual route improvements ranging from 20% to 79% depending on route geometry and wind exposure with only a modest increase in fuel consumption with in the Delhi–Mumbai case study, the schedule-robust solution

requires 330 lb (5.5%) more fuel than the fuel-optimal solution, while still saving 4.7% compared with the conventional ATC route. Intermediate NSGA-II solutions further provide a continuous spectrum of trade-offs between operational efficiency and schedule robustness, allowing dispatchers to select trajectories best aligned with prevailing operational requirements.

Such flexibility is fundamentally absent in conventional single-objective optimisation approaches. A trajectory optimised exclusively for minimum fuel consumption provides no explicit information regarding its sensitivity to forecast uncertainty or its expected schedule reliability. Similarly, methods that minimise expected fuel consumption using probabilistic wind fields still produce only a single "optimal" solution, without quantifying the operational trade-offs associated with alternative routing choices.

The NSGA-II framework addresses this limitation by explicitly resolving the trade-off structure between competing objectives and presenting it in an operationally interpretable form. As a result, the optimisation process becomes not only a tool for identifying efficient trajectories, but also a decision-support framework that enables transparent and informed operational planning under uncertain atmospheric conditions.

7.2 Physical Interpretation of Seasonal Differences

The present study evaluates the framework under the pre-monsoon wind scenario (March 2024), which is characterised by a coherent subtropical westerly jet at approximately 27°N with peak speeds exceeds 50 kt at FL340 (Fig. 2). Under these conditions, NSGA-II routes preferentially exploit the jet core, achieving the fuel savings and robustness improvements reported in Section VI. The monsoon period (July), dominated by the Tropical Easterly Jet at FL310–FL360 and widespread convective activity, would present a contrasting atmospheric environment in which the SIGMET weighting factor W_k in Eq. (10) would become more active. Evaluation of the framework under monsoon conditions is reserved for future work, as the study of seasonal performance sensitivity requires a dedicated set of simulations with monsoon-period GFS data.

7.2 Comparison with Chang et al. (2023)

The arrival-time uncertainty reductions obtained in the present study, ranging from approximately 20% to 79% across the 14 routes with a mean of 49%, are consistent with the 40–65% reduction reported by Chang et al. [5] for the Hong Kong–Amsterdam intercontinental corridor. Similarly, the hypervolume values achieved in this work (mean HV = 0.653) fall within the range of 0.62–0.71 reported in their study. The close agreement between these metrics is significant because it demonstrates that the NSGA-II optimisation framework and the associated trapezoidal collocation robustness formulation remain effective across substantially different operational scales, extending from long-haul oceanic trajectories of approximately 5,700 NM to short- and medium-haul domestic sectors ranging from 145 to 951 NM.

The results also indicate that the compact five-member GFS-based ensemble employed in this study is capable of reproducing the essential characteristics of wind forecast uncertainty relevant for trajectory robustness assessment, despite being substantially smaller than the 51-member ECMWF Ensemble Prediction System used by Chang et al. [5]. This reduction in ensemble size provides considerable computational efficiency while preserving the primary uncertainty structures required for operational optimisation.

The fuel-saving percentages obtained in the present study, ranging from approximately 0.5% to 3.8% of total trip fuel, are lower than the 3–6% savings reported for long-haul intercontinental operations by Chang et al. [5]. Two primary factors contribute to this difference. First, short- and medium-haul

domestic operations are strongly influenced by fixed fuel components, including climb fuel, reserve fuel, descent fuel, alternate fuel, and contingency reserves, which reduce the relative contribution of cruise-phase optimisation to total mission fuel consumption. Second, the simplified five-member "wind stub" ensemble used in this study does not fully reproduce the spatial coherence and intensity of the subtropical jet structures present in operational wind analyses.

Importantly, neither limitation is inherent to the optimisation methodology itself. The use of full operational GFS or ensemble prediction datasets in real-world deployment would be expected to improve the representation of atmospheric flow structures and thereby increase the achievable fuel-saving potential while retaining the demonstrated gains in schedule robustness.

7.4 Limitations of the study

The present study has several limitations. First, the five-member wind ensemble is a simplified representation of operational ensemble weather forecasts. Using larger operational ensembles such as ECMWF EPS would improve uncertainty representation, although at higher computational cost. Second, the chromosome formulation assumes a fixed cruise-altitude structure. Including altitude as an additional optimisation variable could allow the framework to identify more efficient step-climb trajectories automatically. Third, ATC sector-capacity constraints are not included; some trajectories preferred by the schedule-robust solutions may therefore be operationally unavailable under heavy traffic conditions. Fourth, the NASA FLOPS-based performance model carries an estimated accuracy uncertainty of approximately ±3–5%, which affects the reported fuel-saving estimates.

8 Conclusion

This study presents a multi-objective pre-departure trajectory optimisation framework for Indian domestic airspace using the Non-dominated Sorting Genetic Algorithm II (NSGA-II). The framework simultaneously optimises four objectives: fuel consumption, destination arrival-time uncertainty, route-integrated schedule uncertainty, and CO_2 emissions. Instead of producing a single trajectory, the method generates a complete set of optimised flight plans, enabling dispatchers to select trajectories based on operational priorities before departure.

The major findings of the study are summarised as follows:

1. The proposed framework achieves a mean fuel saving of 1.6% relative to the conventional ATC baseline using the simplified five-member wind ensemble. The largest route-specific saving is 3.8% for the VIDP–VOBL (Delhi–Bengaluru) sector.
2. The framework reduces destination arrival-time uncertainty by an average of 49% across all 14 routes. Positive improvements are obtained for every sector, including short-haul routes, demonstrating the effectiveness of uncertainty-aware trajectory optimisation for improving schedule reliability.
3. The NSGA-II solution quality metrics (mean HV = 0.653 and MID = 1.31) are consistent with previously published intercontinental studies despite the shorter domestic route lengths and reduced ensemble size. The C-metric equals 1.0 for all evaluated routes, confirming that at least one NSGA-II solution consistently outperforms the conventional ATC trajectory across all four objectives simultaneously.
4. The optimisation framework remains computationally practical, with execution times of approximately 90–200 seconds per route, making it compatible with operational pre-departure planning timelines without requiring modifications to existing aircraft avionics or ATC procedures.

Applying the observed mean fuel saving of 1.6% to approximately 750 daily Airbus A320neo departures in India, each consuming a mean cruise fuel load of approximately 10,000 lb (4.5 tonnes), yields an estimated annual fleet-level saving of approximately 20,000 tonnes of jet fuel and 63,000 tonnes of CO_2 emissions. At current jet fuel prices of USD 0.60–0.80 per litre, this corresponds to a potential economic benefit of approximately USD 15–20 million annually for the Indian domestic aviation sector.

Overall, the results demonstrate that NSGA-II-based pre-departure trajectory optimisation can provide meaningful improvements in both operational efficiency and schedule robustness within Indian domestic airspace, highlighting its potential as a practical next-generation decision-support tool for airline flight planning under uncertain atmospheric conditions.

Acknowledgement

The authors Acknowledge Padmabhooshan Vasantdada Patil Institute of Technology, Bavdhan, Pune, India, for providing the necessary support, facilities, and academic environment to carry out this research work.

References

1. DGCA, "Annual Report 2023-24," Directorate General of Civil Aviation, Government of India, New Delhi, 2024.
2. B. Ramee, C. Gogu, and S. Doncieux, "Optimal aircraft trajectories for given city pairs using wind forecasts: a multiobjective approach," in Proc. ICRAT, 2020, pp. 1-10.
3. T. Gneiting and A. E. Raftery, "Weather forecasting with ensemble methods," Science, vol. 310, no. 5746, pp. 248-249, Oct. 2005.
4. K. Deb, A. Pratap, S. Agarwal, and T. Meyarivan, "A fast and elitist multiobjective genetic algorithm: NSGA-II," IEEE Trans. Evol. Comput., vol. 6, no. 2, pp. 182-197, Apr. 2002.
5. Z. Chang, M. Hu, and Y. Zhang, "Multi-objective aircraft robust trajectory optimization considering various predictability metrics under uncertain wind," IEEE Access, vol. 11, pp. 104333-104346, 2023. doi: 10.1109/ACCESS.2023.3317285.
6. E. Zermelo, "Uber das Navigationsproblem bei ruhender oder veranderlicher Windverteilung," Z. Angew. Math. Mech., vol. 11, no. 2, pp. 114-124, 1931.
7. C. Schilke and C. Hecker, "Wind-optimal aircraft routing in the lateral and longitudinal plane," in Proc. EIWAC, 2013, pp. 1-8.
8. R. Dancila, R. Botez, and D. Labour, "New flight trajectory optimisation method using genetic algorithms," Aeronautical J., vol. 120, no. 1228, pp. 1-34, 2016.
9. T. Lim and K. Hong, "Stochastic optimal control for aircraft fuel efficiency under wind uncertainty," J. Aerosp. Inf. Syst., vol. 17, no. 4, pp. 210-225, 2020.
10. X. Cheng, H. Chen, and X. Zhang, "Four-dimensional trajectory optimization for green civil aviation using multi-objective evolutionary algorithm," in Proc. IEEE ITSC, 2019, pp. 3482-3487.
11. A. Murrieta-Mendoza, A. Botez, and R. Botez, "Vertical flight trajectory optimization under wind uncertainty using NSGA-II," in Proc. AIAA Aviation Forum, 2019, Paper 2019-3515.
12. R. Gupta and V. Raghunathan, "Wind-optimal routing on Indian domestic routes: a feasibility study," in Proc. CEAS Aerospace Conf., 2022, pp. 1-9.
13. NOAA, "Global Forecast System (GFS) 0.25-degree," National Centers for Environmental Prediction, 2024. [Online]. Available: https://nomads.ncep.noaa.gov

14. Aviation Weather Center, "Convective SIGMET API," NOAA/NWS, 2024. [Online]. Available: https://aviationweather.gov/api
15. Airports Authority of India, "Aeronautical Information Publication (AIP) India," AAI, New Delhi, 2024.
16. NASA, "Flight Optimization System (FLOPS) User's Guide," NASA/TM-2019-220050, 2019.
17. ICAO, "Environmental Technical Manual, Volume II," ICAO Doc 9501, 2015.