

FROM LARGE-SCALE WINDS TO URBAN DECISION MAKING: A CROSS-SCALE FRAMEWORK FOR WIND-AWARE UAV NAVIGATION

Shaoxiang Qin^{1,2} Fuyuan Lyu² Di Zhou³ Xue Liu^{2,4} Xiongye Xiao^{3*}
 Anima Anandkumar^{5*} Liangzhu Leon Wang^{1*}

¹Concordia University ²McGill University ³University of Tennessee, Knoxville
⁴MBZUAI ⁵California Institute of Technology

xxiao9@utk.edu, anima@caltech.edu, leon.wang@concordia.ca

ABSTRACT

Large-scale weather and climate models provide reliable wind information at regional scales, yet their outputs are typically too coarse for direct UAV decision making in geometrically complex urban environments. This paper investigates how large-scale atmospheric information can be transformed into city-scale wind representations and utilized for downstream navigation decisions. We propose a cross-scale prediction and decision framework that takes background wind conditions from existing weather or climate models and combines them with detailed 3D urban geometry to predict time-averaged urban wind fields using a 3D neural operator. The predicted wind fields are then incorporated into a wind-aware UAV trajectory optimization problem to minimize energy consumption under kinematic feasibility and safety constraints. By comparing trajectories planned against a wind-agnostic baseline, we demonstrate significant efficiency gains enabled by AI-predicted wind, specifically 10.3% savings in tailwinds, 7.7% in headwinds, and 3.9% in crosswind conditions. These results indicate that learning decision-relevant urban wind representations offers a practical pathway for bridging large-scale atmospheric information and fine-scale urban decision making.

1 INTRODUCTION

As Unmanned Aerial Vehicles (UAVs) become increasingly prevalent in urban areas, integrating complex local wind fields into autonomous decision making has emerged as a key challenge for improving energy efficiency (Rienecker et al., 2023; Chan et al., 2023; Gu et al., 2025). Although modern climate models provide reliable wind information at regional scales (Kurth et al., 2023; Price et al., 2025), their outputs are typically too coarse to directly support decision making in geometrically complex urban environments, where three-dimensional building structures induce highly heterogeneous flow patterns (Britter & Hanna, 2003). High-fidelity Computational Fluid Dynamics (CFD) can resolve such urban wind fields with high accuracy, but its substantial computational cost (Tominaga et al., 2008) limits its usage. In our setting, steady-state CFD requires around 12 hours per simulation, making CFD impractical for real-time or iterative UAV planning.

To address this gap, we leverage a data-driven model (Qin et al., 2025) based on the 3D Fourier Neural Operator (FNO) (Li et al., 2021; Kovachki et al., 2023) and present a cross-scale framework that connects large-scale atmospheric information with city-scale wind representations suitable for downstream planning. Figure 1 illustrates the overall pipeline. Given background wind conditions that can be obtained from existing weather or climate models, we combine them with detailed 3D urban geometry and predict steady, time-averaged urban wind fields using a model based on 3D FNO. The model is trained on CFD data to map urban geometry to city-scale wind fields under given wind boundary conditions. This approach provides a practical surrogate that preserves key aerodynamic structures induced by urban geometry while enabling fast inference compared to iterative CFD solvers, resulting in over 3 orders of magnitude computational acceleration in our experiments.

*Corresponding authors.

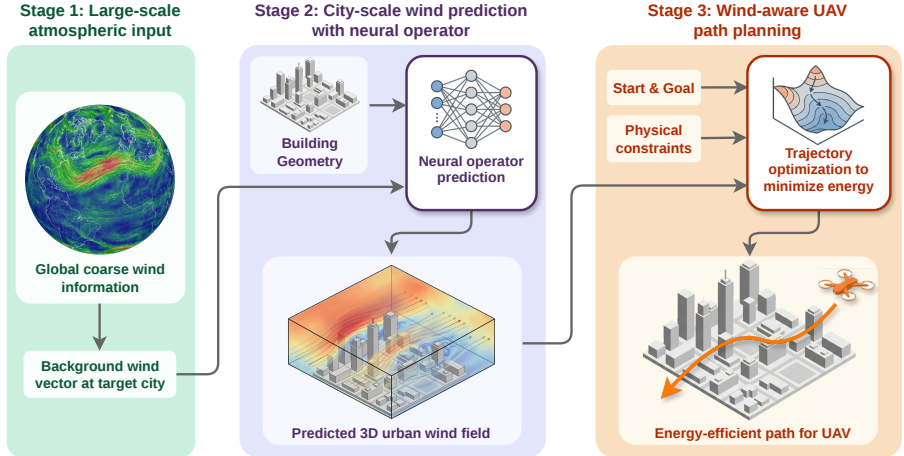


Figure 1: **Overview of the proposed cross-scale framework for urban wind prediction and UAV path planning.** Stage 1 extracts regional background wind conditions from large-scale atmospheric models. Stage 2 predicts city-scale 3D wind fields using urban geometry via a neural operator. Stage 3 generates wind-aware, energy-optimal UAV trajectories under kinematic and safety constraints.

We then incorporate the predicted wind fields into a continuous, wind-aware UAV trajectory optimization pipeline with explicit kinematic feasibility and safety constraints to minimize energy consumption. The planning process follows a two-stage structure. First, a Rapidly-exploring Random Tree (RRT) algorithm (LaValle, 1998) generates multiple collision-free candidate paths to provide diverse feasible initializations in cluttered urban environments. Second, a trajectory optimization (Betts, 1998) routine refines each candidate by jointly optimizing waypoint positions and velocity profiles to minimize total energy consumption under physical constraints. By querying the predicted wind field along the trajectory, the optimizer can exploit favorable airflow patterns and avoid high-drag regions while maintaining safety.

We evaluate the effectiveness of the proposed framework by comparing three planning scenarios: (i) a wind-agnostic baseline assuming zero wind, (ii) planning based on high-fidelity CFD wind fields, and (iii) planning based on FNO-predicted wind fields. Experiments on a held-out urban region demonstrate that the FNO-based approach captures sufficient aerodynamic detail to achieve 10.3% savings in tailwinds, 7.7% in headwinds, and 3.9% in crosswind conditions relative to the baseline. Remarkably, these results recover 90.7%, 77.9%, and 79.6% of the energy savings achieved by planning with high-fidelity CFD, respectively. These results indicate that learning city-scale wind representations from large-scale atmospheric information provides a computationally efficient pathway for bridging climate-scale models and fine-scale urban decision making.

2 DATA OVERVIEW

The 3D neural operator used in this framework was trained on a dataset of 24 time-averaged urban wind fields ($1.2 \text{ km} \times 1.2 \text{ km}$ each) generated in prior research (Qin et al., 2025). These fields were simulated using a GPU-optimized solver CityFFD (Mortezazadeh & Wang, 2019) with large eddy simulation, and the steady-state mean flow was obtained upon convergence. The dataset comprises domains from real-world cities containing high-rise buildings. To validate the generalization of our framework, we perform evaluations on a held-out test city. In the full framework, background wind information is assumed to be provided by large-scale atmospheric models.

3 METHODOLOGY

3.1 URBAN WIND PREDICTION VIA NEURAL OPERATOR

We employ an existing 3D FNO-based (Li et al., 2021; Kovachki et al., 2023) surrogate model (Qin et al., 2025) to infer the wind field. The model utilizes spatial 3D patch-based training with global

geometry encoding. This significantly augments training data, enabling acceptable accuracy from limited full-scale simulations. Specifically, it maps binary building geometries to time-averaged 3D wind fields at a reference speed. Leveraging flow linearity at high Reynolds numbers, these fields can be linearly scaled for other wind speeds. For inference, localized 3D patches are predicted and stitched to reconstruct the full-scale field, ensuring scalability to larger urban environments. This process generates a $1.2 \text{ km} \times 1.2 \text{ km}$ domain in under 30 s, over $1000\times$ faster than CFD. Furthermore, the patch-wise architecture allows for targeted prediction of task-critical regions to further accelerate inference. This efficiency enables practical UAV trajectory planning, whereas traditional CFD requires days, precluding its use in rapid decision making.

3.2 WIND-AWARE UAV PATH PLANNING

The planning module computes energy-optimal trajectories using a two-stage approach: sampling-based initialization followed by continuous trajectory optimization.

Initialization via RRT. A Rapidly-exploring Random Tree (LaValle, 1998) first generates multiple topologically distinct, collision-free paths. These paths provide diverse warm starts to alleviate local minima issues in the subsequent optimization. Details are provided in Appendix B.2.

Continuous trajectory optimization. We refine each candidate path as a nonlinear program (NLP) solved via the Interior Point OPTimizer (IPOPT) (Wächter & Biegler, 2006), after which we select the trajectory with the minimum energy cost as the final mission plan. Decision variables include the discretized UAV states, consisting of position \mathbf{x}_i , velocity \mathbf{v}_i , and acceleration \mathbf{a}_i at each waypoint, and a uniform time step Δt . The NLP is defined by the following components: **(i) Objective function:** We minimize the time integral of the UAV power as its total energy consumption, where the power model (Appendix A) accounts for useful, induced, and profile power. The useful power term accounts for the work done against aerodynamic drag and gravity. This term implicitly drives the solver to exploit tailwinds and avoid high-drag zones. Induced and profile power components model the aerodynamic losses of the propellers. **(ii) Kinematic feasibility:** Trapezoidal collocation constraints enforce physical consistency between position, velocity, and acceleration across time steps. **(iii) Geometric feasibility:** A hard safety constraint ensures the distance to the nearest obstacle exceeds a predefined clearance at all nodes. **(iv) Actuation limits:** Bounds are imposed on velocity, acceleration, and thrust magnitudes to respect the UAV’s operational envelope. See Appendix B.3 for details of ii–iv.

4 RESULTS AND DISCUSSION

Experimental Setup. We evaluate the framework using a $1.2 \text{ km} \times 1.2 \text{ km}$ test urban area excluded from the FNO training dataset. Background inflow is characterized by a reference speed of 4 m/s at a 10 m height. We sample 256 random pairs of start and goal locations across the domain, with both terminal points set at 50 m altitude and the allowable flight envelope ranging from 30 m to 150 m.

For each mission, we compare three planning strategies: (i) a wind-agnostic baseline assuming zero wind velocity, (ii) planning with high-fidelity wind information from CFD simulations, and (iii) planning with wind estimates from the FNO-based model. Strategy (ii) serves as a performance upper bound, as the prohibitive computational latency of CFD precludes its use for practical mission planning. To ensure a robust comparison and mitigate the effect of local minima, each strategy employs 10 distinct RRT initializations per mission. The final energy consumption for all three methods is evaluated within the high-fidelity CFD wind field to reflect the actual physical conditions the UAV would encounter during operation.

4.1 WIND FIELD PREDICTION AND TRAJECTORY CHARACTERISTICS

The CFD-simulated wind field and the FNO-based prediction are compared in the first column of Figure 2. The FNO model achieves a mean velocity error of approximately 0.6 m/s while requiring less than one-thousandth of the computational time of full CFD. Although minor fine-scale discrepancies remain, the model accurately captures dominant aerodynamic structures, including high-velocity zones in urban canyons and recirculation regions behind buildings.

These wind representations directly influence trajectory planning. Wind-agnostic paths favor near-straight shortest routes, whereas wind-aware planners generate curvilinear trajectories that exploit local flow structures and vertical wind gradients. Specifically, the UAV climbs in tailwinds to harness stronger assisting flows and descends in headwinds to reduce resistance. Figure 2 shows the urban height map and the optimized trajectories under headwind, tailwind, and crosswind conditions. Trajectories computed using the FNO-based model closely match those obtained with high-fidelity CFD. Despite approximation errors, our approach consistently identifies more energy-efficient routes compared to the wind-agnostic baseline.

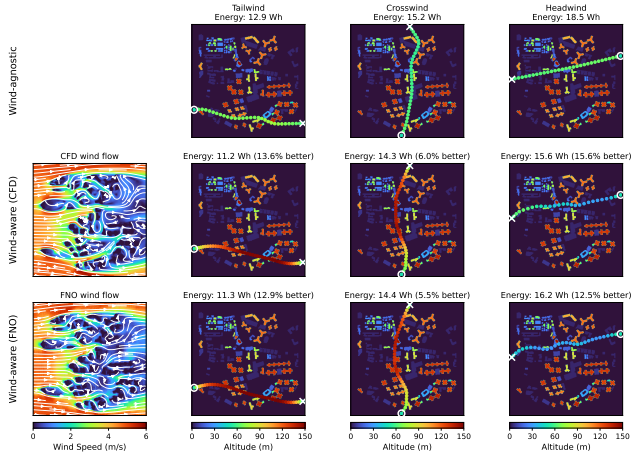


Figure 2: Urban wind fields and optimized trajectories comparing wind-agnostic, CFD-based, and FNO-based planning under tailwind, crosswind, and headwind conditions. Circles for start, crosses for goal.

4.2 STATISTICAL ANALYSIS OF TRAJECTORY CHARACTERISTICS

Figure 3 displays individual mission results in polar coordinates, where the angular dimension represents the relative wind angle θ , defined between the start-to-goal mission vector and the global wind direction. For quantitative analysis, the missions are categorized into: (i) Tailwind ($|\theta| \leq 45^\circ$), (ii) Headwind ($|\theta - 180^\circ| \leq 45^\circ$), and (iii) Crosswind (all other angles).

The left subplot illustrates energy savings relative to the wind-agnostic baseline, highlighting that savings are most pronounced in tailwind and headwind missions while being relatively lower in crosswind conditions. FNO-based trajectories recover 90.7% of the energy savings attained using high-fidelity CFD in tailwinds, 77.9% in headwinds, and 79.6% in crosswinds. Variance in savings is notably greater in headwinds. While the tailwind strategy, climbing to exploit wind gradients at altitude, remains uniform, the headwind strategy requires navigating complex flow structures induced by building geometries, rendering efficiency highly dependent on local urban topology.




Figure 3: Trajectory performance statistics across 256 missions as a function of relative wind angle, including energy savings, mean flight altitude, and mean ground speed.

The middle subplot shows mean flight altitude correlating with the relative wind angle. To exploit the vertical wind gradient, FNO-based trajectories climb in tailwinds (mean 125.4 m vs. 69.9 m baseline) and descend in headwinds (60.0 m vs. 70.9 m baseline) to minimize drag. The wind-agnostic baseline fails to exploit these vertical structures, maintaining a nearly constant altitude across all directions.

The right subplot displays ground speed distributions. The wind-agnostic baseline remains at 13.5 m/s, whereas FNO-based planning leverages the flow to reach 16.3 m/s in tailwinds and reduces to 13.0 m/s in headwinds to maintain efficiency. The close alignment with high-fidelity CFD results suggests that FNO-inferred wind fields provide a viable basis for practical, wind-aware path planning in urban environments.

REFERENCES

- John T. Betts. Survey of numerical methods for trajectory optimization. *Journal of Guidance, Control, and Dynamics*, 21(2):193–207, 1998.
- Rex Edward Britter and Steven R. Hanna. Flow and dispersion in urban areas. *Annual Review of Fluid Mechanics*, 35(1):469–496, 2003.
- Y.Y. Chan, Kam K.H. Ng, C.K.M. Lee, Li-Ta Hsu, and K.L. Keung. Wind dynamic and energy-efficiency path planning for unmanned aerial vehicles in the lower-level airspace and urban air mobility context. *Sustainable Energy Technologies and Assessments*, 57:103202, 2023.
- Ruijia Gu, Yifei Zhao, and Xinhui Ren. Integrating wind field analysis in UAV path planning: Enhancing safety and energy efficiency for urban logistics. *Chinese Journal of Aeronautics*, pp. 103605, 2025.
- Nikola Kovachki, Zongyi Li, Burigede Liu, Kamyar Azizzadenesheli, Kaushik Bhattacharya, Andrew Stuart, and Anima Anandkumar. Neural operator: Learning maps between function spaces with applications to PDEs. *Journal of Machine Learning Research*, 24(89):1–97, 2023.
- Thorsten Kurth, Shashank Subramanian, Peter Harrington, Jaideep Pathak, Morteza Mardani, David Hall, Andrea Miele, Karthik Kashinath, and Anima Anandkumar. Fourcastnet: Accelerating global high-resolution weather forecasting using adaptive fourier neural operators. In *Proceedings of the Platform for Advanced Scientific Computing Conference*, pp. 1–11, 2023.
- Steven M. LaValle. Rapidly-exploring random trees: A new tool for path planning. Technical report, Iowa State University, 1998.
- Zongyi Li, Nikola Borislavov Kovachki, Kamyar Azizzadenesheli, Burigede Liu, Kaushik Bhattacharya, Andrew M. Stuart, and Anima Anandkumar. Fourier neural operator for parametric partial differential equations. In *9th International Conference on Learning Representations, ICLR 2021, Virtual Event, Austria, 2021*. OpenReview.net. URL <https://openreview.net/forum?id=c8P9NQVtmnO>.
- Mohammad Mortezaazadeh and Liangzhu Wang. An adaptive time-stepping semi-Lagrangian method for incompressible flows. *Numerical Heat Transfer, Part B: Fundamentals*, 75(1):1–18, 2019.
- Ilan Price, Alvaro Sanchez-Gonzalez, Ferran Alet, Tom R Andersson, Andrew El-Kadi, Dominic Masters, Timo Ewalds, Jacklynn Stott, Shakir Mohamed, Peter Battaglia, et al. Probabilistic weather forecasting with machine learning. *Nature*, 637(8044):84–90, 2025.
- Shaoliang Qin, Dongxue Zhan, Ahmed Marey, Dingyang Geng, Theodore Potsis, and Liangzhu Leon Wang. Data-efficient rapid prediction of urban airflow and temperature fields for complex building geometries. *arXiv preprint arXiv:2503.19708*, 2025.
- Hannes Rienecker, Veit Hildebrand, and Harald Pfifer. Energy optimal 3D flight path planning for unmanned aerial vehicle in urban environments. *CEAS Aeronautical Journal*, 14(3):621–636, 2023.
- Yoshihide Tominaga, Akashi Mochida, Ryuichiro Yoshie, Hiroto Kataoka, Tsuyoshi Nozu, Masaru Yoshikawa, and Taichi Shirasawa. AIJ guidelines for practical applications of CFD to pedestrian wind environment around buildings. *Journal of Wind Engineering and Industrial Aerodynamics*, 96(10-11):1749–1761, 2008.
- Andreas Wächter and Lorenz T Biegler. On the implementation of an interior-point filter line-search algorithm for large-scale nonlinear programming. *Mathematical Programming*, 106(1): 25–57, 2006.

A ENERGY CONSUMPTION MODEL

Urban wind fields significantly alter the aerodynamic loading of a UAV and, consequently, its energy consumption. To enable wind-aware trajectory optimization with tractable computational cost, we adopt a quasi-steady power model evaluated along the discretized trajectory. This model explicitly accounts for spatially varying wind fields through the air-relative velocity vector $\mathbf{v}_{a,i}$ at each node. Specifically, the mechanical power is decomposed into three fundamental components: useful power, induced power, and profile power. This formulation allows the optimizer to account for the work done against aerodynamic drag and gravity while modeling the internal losses of the propulsion system.

A.1 STATE, WIND FIELD, AND BASIC DEFINITIONS

We consider a trajectory sampled at N discrete nodes indexed by $i \in \{0, 1, \dots, N-1\}$. At node i , the UAV state is

$$\mathbf{x}_i \in \mathbb{R}^3, \quad \mathbf{v}_i \in \mathbb{R}^3, \quad \mathbf{a}_i \in \mathbb{R}^3, \quad (1)$$

representing position, ground-relative velocity, and ground-relative acceleration, respectively. The (predicted) wind field is a vector function

$$\mathbf{w}(\mathbf{x}) \in \mathbb{R}^3, \quad (2)$$

which returns the local wind velocity at a spatial location \mathbf{x} .

The air-relative velocity (airspeed vector) at node i is defined as

$$\mathbf{v}_{a,i} = \mathbf{v}_i - \mathbf{w}(\mathbf{x}_i), \quad (3)$$

with magnitude $\|\mathbf{v}_{a,i}\|$. Let $\mathbf{g} \in \mathbb{R}^3$ denote the gravitational acceleration vector (pointing downward), with magnitude $g = \|\mathbf{g}\|$.

A.2 AERODYNAMIC DRAG AND REQUIRED THRUST

We model the parasite drag force using a quadratic form:

$$\mathbf{D}_i = -\frac{1}{2}\rho C_d A_f \|\mathbf{v}_{a,i}\| \mathbf{v}_{a,i}, \quad (4)$$

where ρ is air density, C_d is an effective drag coefficient, and A_f is the UAV frontal reference area. The negative sign indicates that drag opposes the air-relative motion.

Given \mathbf{a}_i and the drag force, the thrust vector required to satisfy translational dynamics is

$$\mathbf{T}_i = m\mathbf{a}_i - m\mathbf{g} - \mathbf{D}_i, \quad (5)$$

where m is the UAV mass. We denote the thrust magnitude and the unit thrust direction by

$$T_i = |\mathbf{T}_i|, \quad \hat{\mathbf{t}}_i = \frac{\mathbf{T}_i}{T_i}. \quad (6)$$

A.3 POWER AT EACH NODE

Following standard rotorcraft momentum theory, we define the component of air-relative velocity along the thrust axis as

$$v_{c,i} = \mathbf{v}_{a,i}^\top \hat{\mathbf{t}}_i. \quad (7)$$

The induced velocity v_i^{ind} is computed by

$$v_i^{\text{ind}} = -\frac{1}{2}v_{c,i} + \sqrt{\left(\frac{1}{2}v_{c,i}\right)^2 + \frac{T_i}{2\rho A}}, \quad (8)$$

where A is the total rotor disk area.

We then compute three power components at node i :

$$P_{u,i} = \mathbf{T}_i^\top \mathbf{v}_{a,i}, \quad (9)$$

$$P_{ind,i} = T_i v_i^{ind}, \quad (10)$$

$$P_{p,i} = P_{p,hover} \left(\frac{T_i}{mg} \right)^{\frac{3}{2}}, \quad (11)$$

where $P_{u,i}$ is the useful power, representing the rate of work done by the thrust vector \mathbf{T}_i against the air-relative velocity $\mathbf{v}_{a,i}$. This term accounts for the power required to overcome parasite drag and to facilitate changes in the UAV’s mechanical energy, including both potential and kinetic energy variations. The induced power $P_{ind,i}$ accounts for the energy required to accelerate air through the rotor disks to generate thrust T_i . The profile power $P_{p,i}$ models the aerodynamic losses due to the rotation of the rotor blades, which is scaled from the hover profile power $P_{p,hover}$ based on the thrust magnitude relative to the UAV weight mg .

The total power at node i is the sum of these components:

$$P_i = P_{u,i} + P_{ind,i} + P_{p,i}. \quad (12)$$

A.4 TRAJECTORY ENERGY INTEGRATION

Given a set of time intervals $\{\Delta t_i\}_{i=0}^{N-2}$ between consecutive nodes, the total energy is obtained by trapezoidal integration:

$$E = \sum_{i=0}^{N-2} \frac{1}{2} (P_i + P_{i+1}) \Delta t_i. \quad (13)$$

In our trajectory optimization formulation, we impose a uniform time step $\Delta t_i \equiv \Delta t$ as a constraint to couple the discretized kinematics and the energy objective. For post-optimization evaluation, the actual time interval Δt_i is derived from the optimized positions and velocities, i.e., $\Delta t_i \approx \|\mathbf{x}_{i+1} - \mathbf{x}_i\| / (\frac{1}{2}(\|\mathbf{v}_i\| + \|\mathbf{v}_{i+1}\|))$, ensuring that the energy consumption reflects the actual recovered velocity profile.

A.5 MODEL PARAMETERS

Unless otherwise stated, the parameters in equation 4–equation 11 are treated as constants: ρ (air density), C_d (drag coefficient), A_f (frontal area), m (mass), A (rotor disk area), g (gravitational acceleration magnitude), and $P_{p,hover}$ (hover profile power). The wind field $\mathbf{w}(\mathbf{x})$ is provided by the learned wind predictor and queried continuously along the trajectory, enabling the planner to exploit favorable airflow structures while maintaining physical consistency.

B ENERGY-AWARE PATH PLANNING VIA RRT INITIALIZATION AND CONTINUOUS TRAJECTORY OPTIMIZATION

This work formulates wind-aware path planning in urban wind fields as a two-stage pipeline that balances feasibility, computational efficiency, and physical consistency. The planner first employs sampling-based initialization to rapidly generate collision-free candidate paths in the 3D environment. These paths are then refined by a continuous trajectory optimization routine that explicitly minimizes energy consumption under kinematic feasibility and safety constraints.

B.1 PROBLEM SETTING AND INPUTS

The environment is represented by a 3D occupancy grid and an associated signed distance field (SDF). The occupancy grid encodes obstacle voxels, while the SDF provides a continuous measure of distance to the nearest obstacle, enabling a clearance margin to be enforced both in discrete planning and continuous optimization. A spatially varying wind field is provided as a 3D velocity vector function $\mathbf{w}(\mathbf{x})$, obtained from CFD simulations or the learned wind predictor.

Given a start location \mathbf{x}_s and a goal location \mathbf{x}_g , the objective is to compute a trajectory that is (i) collision-free with a prescribed clearance, (ii) dynamically feasible under realistic limits on velocity, acceleration, and thrust, and (iii) energy-minimal in the presence of wind.

B.2 FEASIBLE PATH INITIALIZATION WITH RRT

To obtain an initial feasible solution quickly, we employ a Rapidly-exploring Random Tree (RRT) planner in the 3D grid. RRT incrementally grows a tree rooted at the start by repeatedly sampling random points in free space and connecting each sample to the nearest existing node via a short forward extension. This approach is well-suited for cluttered urban scenes because it explores large free-space volumes efficiently and does not require a global search over all grid cells.

In our implementation, the extension length is adapted using the local SDF value: when the sampled region is far from obstacles, larger steps are taken to accelerate exploration; when near obstacles, the step size is reduced to increase maneuvering resolution and avoid collisions. Each newly proposed tree edge is validated by discretizing the line segment and checking obstacle clearance conditions. The planner terminates as soon as the tree reaches the goal region and returns a piecewise-linear waypoint sequence that is guaranteed to satisfy the geometric obstacle constraints at the chosen checking resolution.

Because RRT is stochastic, repeated runs naturally produce diverse feasible paths with different homotopy classes and different exposure to wind structures. We therefore generate multiple candidate initial paths and refine each candidate in the next stage, retaining the best-performing solutions.

B.3 CONTINUOUS TRAJECTORY OPTIMIZATION FOR ENERGY MINIMIZATION

While the RRT output is collision-free, it is not optimized for energy and does not explicitly enforce dynamic limits. We therefore refine the waypoint sequence using a continuous trajectory optimization formulation that jointly adjusts the waypoint positions and the associated velocity profile.

Let the trajectory be discretized into N nodes indexed by $i \in \{0, \dots, N-1\}$. The decision variables are

$$\mathcal{Z} = \{\mathbf{x}_i, \mathbf{v}_i, \mathbf{a}_i\}_{i=0}^{N-1} \cup \{\Delta t\}, \quad (14)$$

where $\mathbf{x}_i \in \mathbb{R}^3$ is the position, $\mathbf{v}_i \in \mathbb{R}^3$ is the ground-relative velocity, $\mathbf{a}_i \in \mathbb{R}^3$ is the acceleration, and Δt is a uniform time step shared across segments. The wind field enters through the air-relative velocity $\mathbf{v}_{a,i} = \mathbf{v}_i - \mathbf{w}(\mathbf{x}_i)$, which affects both drag and the required thrust, and consequently the power and energy.

Objective. Using the per-node power model defined in Section A, the trajectory optimization minimizes the total energy via trapezoidal integration:

$$\min_{\mathcal{Z}} E = \sum_{i=0}^{N-2} \frac{1}{2} (P_i + P_{i+1}) \Delta t. \quad (15)$$

Kinematic consistency. To ensure kinematic consistency between position, velocity, and acceleration, we impose a trapezoidal discretization of the continuous-time dynamics:

$$\mathbf{x}_{i+1} = \mathbf{x}_i + \frac{1}{2} (\mathbf{v}_i + \mathbf{v}_{i+1}) \Delta t, \quad (16)$$

$$\mathbf{v}_{i+1} = \mathbf{v}_i + \frac{1}{2} (\mathbf{a}_i + \mathbf{a}_{i+1}) \Delta t, \quad (17)$$

for all $i = 0, \dots, N-2$, together with boundary conditions

$$\mathbf{x}_0 = \mathbf{x}_s, \quad \mathbf{x}_{N-1} = \mathbf{x}_g. \quad (18)$$

Safety constraints (collision avoidance with clearance). We enforce obstacle clearance using the SDF:

$$\text{SDF}(\mathbf{x}_i) \geq d_{\text{clr}}, \quad \forall i, \quad (19)$$

where d_{clr} is the required clearance. To reduce the risk of cutting corners between nodes, we additionally check a small number of interior samples along each segment. Let

$$\mathbf{x}_{i,k} = (1 - \alpha_k) \mathbf{x}_i + \alpha_k \mathbf{x}_{i+1}, \quad \alpha_k \in (0, 1), \quad (20)$$

then we impose

$$\text{SDF}(\mathbf{x}_{i,k}) \geq d_{\text{clr}}, \quad \forall i, \forall k, \quad (21)$$

which empirically provides a strong safety margin while keeping computation efficient.

Dynamic limits. We impose pointwise bounds on airspeed, acceleration, and thrust magnitude:

$$\|\mathbf{v}_{a,i}\| \leq v_{\max}, \quad \|\mathbf{a}_i\| \leq a_{\max}, \quad \|\mathbf{T}_i\| \leq T_{\max}, \quad \forall i, \quad (22)$$

where $\mathbf{v}_{a,i} = \mathbf{v}_i - \mathbf{w}(\mathbf{x}_i)$ is the air-relative velocity (airspeed) accounting for the local wind field $\mathbf{w}(\mathbf{x}_i)$. The thrust vector \mathbf{T}_i is determined by the dynamics and aerodynamics defined in Section A. These constraints ensure that the optimized trajectory is not only collision-free but also executable by the UAV within its physical actuator capability.

B.4 NUMERICAL SOLUTION

The resulting problem is a nonlinear program (NLP) with a nonlinear objective and constraints due to the wind-coupled aerodynamics and the SDF-based collision constraints. We solve this NLP using a gradient-based interior-point solver IPOPT. In practice, the quality of the solution depends on initialization because the optimization landscape is non-convex. We therefore adopt a multi-start strategy: multiple RRT runs provide diverse feasible initial paths, each of which is refined by the trajectory optimizer. The final output is selected as the lowest-energy solution among the optimized candidates. This procedure yields robust performance in complex urban scenes while enabling the planner to systematically exploit wind structures to reduce energy expenditure.

B.5 WIND-AWARE VS. WIND-AGNOSTIC BASELINES

To quantify the benefit of incorporating predicted wind, we report two variants of the pipeline:

- **Wind-aware planning:** the optimizer uses $\mathbf{w}(\mathbf{x})$ when computing air-relative velocity, drag, thrust, and power.
- **Wind-agnostic planning:** the optimizer assumes $\mathbf{w}(\mathbf{x}) \equiv \mathbf{0}$ during optimization.

For fair comparison, both solutions are evaluated under the same wind field during energy computation.

C STATISTICAL CHARACTERISTICS OF TRAJECTORIES

The table below presents the specific statistical values for the data shown in Figure 3.

Table 1: Quantitative comparison of flight characteristics under different wind conditions. Results are averaged over 256 test missions. Base: wind-agnostic baseline; CFD: high-fidelity CFD; FNO: AI-predicted wind. All energy savings are calculated relative to the wind-agnostic baseline.

Condition	Energy Savings (%)		Mean Altitude (m)			Mean Speed (m/s)		
	CFD	FNO	Base	CFD	FNO	Base	CFD	FNO
Tailwind	11.4	10.3	69.9	124.5	125.4	13.5	16.2	16.3
Headwind	9.9	7.7	70.9	61.1	60.0	13.5	12.9	13.0
Crosswind	4.9	3.9	70.0	80.1	82.6	13.5	13.8	13.8

D MODEL ARCHITECTURE

The 3D neural operator architecture is based on the existing FNO framework, where each block consists of a skip connection in parallel with a Fourier layer followed by a feedforward network. To enhance feature representation, the feedforward component utilizes a gated MLP architecture rather than a standard MLP. A critical design for accurate wind field prediction is the global geometry encoding, which is computed before the spatial 3D patch-wise training. This encoding utilizes multi-directional distance features that represent the distance from each spatial point to the nearest obstacle across a dense set of angles. Compared to a standard signed distance function, these features explicitly capture obstacle orientation and bearing, providing a more informative geometric representation that facilitates the learning of direction-dependent structures. More details can be found in the literature (Qin et al., 2025).