Energy-Efficient Trajectory Planning and Feasibility Assessment Framework for Drone Package Delivery *

Abenezer G. Taye† and Peng Wei‡
George Washington University, Washington, DC, 20052, USA

This paper presents a comprehensive framework for energy-efficient trajectory planning and feasibility assessment in drone package delivery operations across urban areas. The framework is designed to assist unmanned aerial systems (UAS) operators within the Unmanned Aircraft System Traffic Management (UTM) architecture by generating feasible operational plans for a given set of package delivery missions. The framework addresses the multifaceted computational challenges of this problem using a two-layer approach. The upper layer consists of a Markov Decision Process (MDP)-based trajectory planner that generates energy-efficient trajectories by considering factors such as the aircraft model, the length and smoothness of the trajectories, and the interaction of the aircraft with the wind field. The lower layer then assesses the feasibility of these energy-efficient trajectories through a battery state of charge (SoC) prediction-based uncertainty quantification scheme. We demonstrate the efficacy of the proposed framework in a realistic drone package delivery scenario set in an urban environment in Boston City. In addition, an actual flight experiment was conducted to validate the framework's performance. The results highlight the framework's capability to generate energy-efficient trajectories and assess the feasibility of each mission. Source code: `github.com/Abenezergirma/ETP-FA-Package-Delivery`

I. Introduction

A. Motivation

Drone package delivery using small unmanned aerial systems (sUAS) is experiencing accelerated progress and is closer to becoming reality. The FAA recently granted approval for companies such as Zipline and Wing Aviation to operate commercial drones in the Dallas-Fort Worth (DFW) area without requiring visual observers, thus enabling beyond-visual line of sight (BVLOS) operations [1–3]. This historic authorization marked a significant shift in the regulatory landscape, aiming to incorporate routine drone deliveries safely into the national airspace. It allows leading

*This work was presented in part at the AIAA SCITECH 2024 Forum, Taye, A. G., and Wei, P. (2024) "Flight Mission Feasibility Assessment of Urban Air Mobility Operations under Battery Energy Constraint," AIAA Paper 2024-0532, Orlando, FL, January 2024.

†PhD Candidate, Department of Mechanical & Aerospace Engineering, abenezertaye@gwu.edu, AIAA Student Member.
‡Associate Professor, Department of Mechanical & Aerospace Engineering, pwei@gwu.edu, AIAA Associate Fellow.

drone delivery companies to deliver packages while keeping their drones safely separated using Unmanned Aircraft System Traffic Management (UTM) technology [2].

The UTM Concept of Operations (ConOps) proposed by the FAA [4] serves as a blueprint outlining how sUAS operations are intended to function. It consists of an architecture with various components, among which sUAS operators play a critical role. The primary responsibility of sUAS operators is to generate operational plans, which are developed "prior to the operation and indicate a four-dimensional (4D) volume of airspace within which the operation is expected to occur" [4]. After these plans are generated, they must be submitted to sUAS service suppliers (USS) for final flight authorization. Therefore, the framework proposed in this paper can be viewed as an automated tool that assists sUAS operators in generating these operational plans.

The framework proposed in this paper specifically addresses the challenge of generating energy-efficient trajectories while ensuring the safe operation of the aircraft involved. In this context, the safety of the sUAS system is considered from two perspectives. The first aspect is collision avoidance between participating aircraft, which is achieved through separation assurance using a reachability analysis-based collision avoidance scheme. The second aspect addresses the risk associated with insufficient battery energy, as sUAS operations typically rely on electric-powered unmanned aerial vehicles (UAVs). To mitigate this risk, the framework generates energy-efficient trajectories and conducts battery energy-related feasibility assessments on these trajectories. Figure 1 illustrates a drone package delivery scenario in an urban environment, demonstrating how the framework generates safe, energy-efficient flight trajectories and performs battery-related feasibility assessments.

B. Contributions

This work presents several key contributions to the field of sUAS trajectory planning and battery energy-related feasibility assessment for sUAS operations. These contributions are outlined as follows:

1) **Energy-efficient sUAS trajectory planning in the presence of wind:** We developed a Markov Decision Process-based trajectory planner that generates energy-efficient flight trajectories, accounting for various aircraft and environmental factors, including wind field, the length and smoothness of trajectories, that influence the energy consumption of the flight mission.
2) **Battery energy-related flight mission feasibility assessment:** A feasibility assessment scheme based on battery state of charge (SoC) prediction was developed to account for uncertainties in battery states. This scheme evaluates the feasibility of the generated trajectories and determines whether to proceed with the flight mission.
3) **Realistic drone package delivery scenario:** A realistic drone package delivery scenario was developed for a region in Boston City, incorporating factors such as a wind field in the urban environment, to demonstrate the practical application of the proposed framework in a real-world setting.
4) **Flight test validation:** An actual flight experiment was conducted to validate the performance of the feasibility

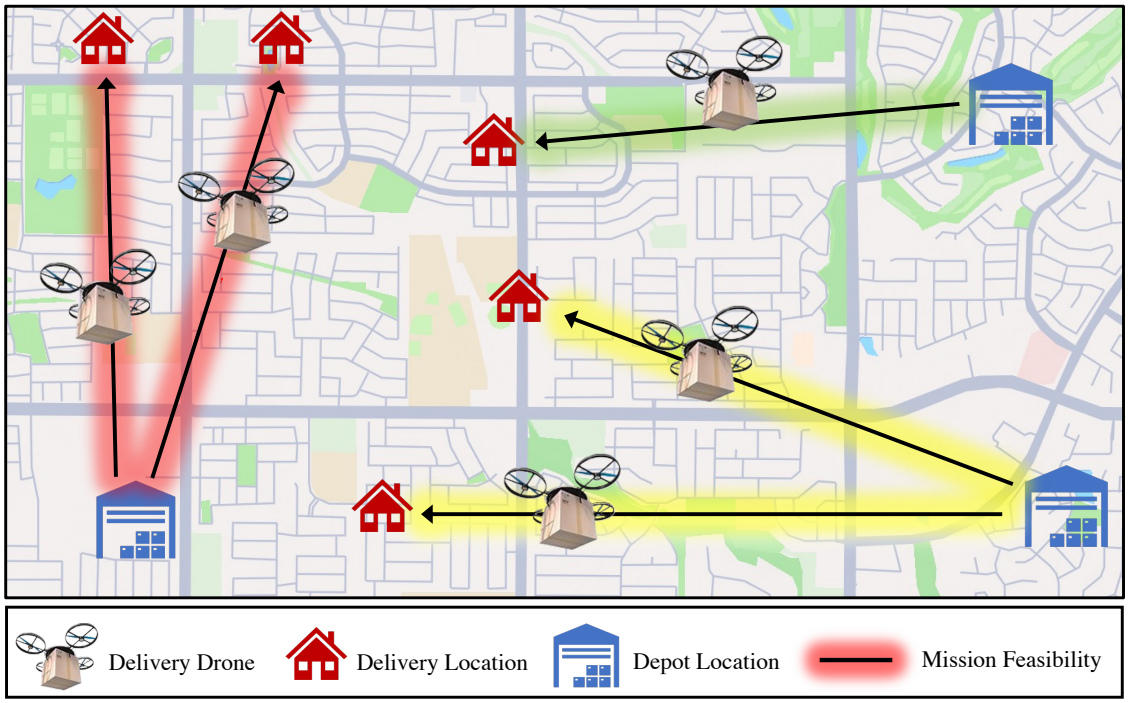

Fig. 1 Our proposed framework generates safe and energy-efficient trajectories for each aircraft in the environment, followed by battery energy-related feasibility assessments for the generated trajectories.

assessment scheme using real flight data.

The rest of the paper is structured as follows. Section II highlights existing research relevant to our work. Section III provides a brief discussion of the necessary background material related to the proposed framework. In Section IV, we present the problem formulation and give an overview of our two-layer approach. We then elaborate on the components of the upper-layer trajectory planner in Section V and the reward functions in Section VI. Section VII discusses the lower feasibility assessment layer, while Section VIII explains the overall framework's working procedure. Finally, in Section IX, we present the package delivery scenario developed for an area near Boston City, including experimental results and hardware tests conducted to validate the feasibility assessment. The paper concludes with Section X.

II. Related Work

We now summarize three areas of previous work related to our overall problem: multi-agent trajectory planning, trajectory optimization, and flight mission feasibility assessment. They utilize a range of concepts from artificial intelligence, operations research, and prognostics.

A. Multi-Agent Trajectory Planning

The literature on multi-agent trajectory planning algorithms is extensive and can be broadly classified into centralized and decentralized methods. Centralized methods involve a central controller that, using data from sensors and radar,

observes the state of each aircraft, obstacles, and trajectory constraints to precompute trajectories for all aircraft, typically within an optimal control framework. Various techniques, such as semidefinite programming [5], nonlinear programming [6, 7], and mixed-integer linear programming [8–11] are employed to resolve conflicts and pursue a global optimum. In addition, evolutionary techniques [12], reinforcement learning [13, 14], and particle swarm optimization [15] based approaches have been proposed. However, these methods often face scalability issues as the number of aircraft and obstacles increases and require recomputation when new aircraft enter the scene.

In contrast, decentralized methods, where each aircraft resolves conflicts locally in either a cooperative or non-cooperative manner, scale better with the number of aircraft and objects in the system, although they may not achieve globally optimal solutions. Decentralized approaches are often more robust, avoiding a single point of failure, but involve trade-offs between computational scalability and solution quality. For instance, decentralized trajectory planning frameworks such as MADER ensure collision-free paths in dynamic 3D environments through asynchronous trajectory optimization [16], while spatiotemporal occupancy grid maps have been employed to enhance real-time collision avoidance for multiple UAVs in dynamic settings [17]. Similarly, probabilistic safety guarantees under uncertainty can be achieved through reinforcement learning combined with robust optimization, as demonstrated in [18]. Further, distributed model predictive control has shown scalability in handling multi-agent collision avoidance [19]. Building on this, we previously proposed a Markov Decision Process (MDP)-based decentralized trajectory planning algorithm that is highly scalable, operates in a 3D environment, and is formally verified [20]. We now extend this trajectory planner to generate energy-efficient trajectories in addition to ensuring the safe operation of the aircraft.

B. UAV Trajectory Optimization in the Presence of Wind

An essential component of the trajectory planning framework presented in this paper is the generation of energy-efficient trajectories under wind influence, commonly referred to as trajectory optimization in the literature [21]. Various approaches have been proposed, including customized versions of the A-Star algorithm, such as those by Rienecker et al. [22] and Ebert et al. [23], which enable UAVs to exploit local wind patterns. Hong et al. [13] introduced a Twin-Delayed Deep Deterministic Policy Gradient (TD3) algorithm for real-time, energy-efficient path planning in windy conditions, while Baskar et al. [24] developed a probabilistic roadmap (PRM) path planning framework that dynamically balances energy and navigation costs in response to wind. Additionally, Al et al. [25] proposed an MDP-based approach prioritizing actions that align the UAV's path with favorable wind components, and Banerjee et al. [26] introduced a Deep Q-learning-based method that incorporates a surrogate battery model to penalize voltage drop in energy-intensive trajectories. Souto et al. [27] presented a Q-learning approach adjusting UAV altitude to minimize wind exposure by maintaining altitude above the tallest buildings in each grid cell.

However, most of these approaches rely on simplified models with two-dimensional discrete state spaces, limiting their practical application. Moreover, they generally focus on wind interaction at a single decision-making step, without

fully accounting for factors that impact overall trajectory energy consumption, such as the aircraft model, powertrain characteristics, and trajectory length and smoothness. To address these limitations, we propose a method that generates energy-efficient trajectories by integrating all these critical factors.

C. Flight Mission Feasibility Assessment

In the lower layer of our framework, we conduct battery energy-related flight mission feasibility assessment. Recent works in this area include Corbetta et al. [28] presented a framework to quantify uncertainty in mission success due to available battery energy, but only considered the power train model of a UAV. Quinones-Grueiro et al. [29] studied the UAV navigation problem in urban environments while taking the battery energy into account. Additionally, Schumann et al. [30] introduced a prognostics-as-a-service (PaaS) framework for UAVs that monitors the health and state of charge of the battery (SoC), among other features of the UAV online, communicates the state of the aircraft to ground controllers, and can perform contingency planning. Furthermore, in [31, 32], we considered complex aircraft and battery dynamic models, along with actual wind forecasts, to examine the impact of battery energy on flight missions.

III. Background

This section provides the necessary background material to understand the trajectory planning and feasibility assessment layers of the proposed framework. It begins with a brief discussion of the fundamentals of Markov Decision Processes (MDPs), followed by an overview of the prognostics architecture used to assess the feasibility of the flight mission. Finally, the section formulates the entire trajectory planning and feasibility assessment framework.

A. Markov Decision Process

An MDP is defined as a tuple (S, A, R, T), where S is the set of states, A is the set of actions, $R(s, a)$ is the reward function that assigns a reward for taking action $a \in A$ from state $s \in S$, and $T(s, a, s')$ is the transition function that describes the dynamics of the environment. The solution to the MDP, the optimal policy π^*, specifies the optimal action $a^* \in A$ to take from each state $s \in S$ to maximize the expected return. The maximum expected value obtained from each state $s \in S$ is represented by the optimal value function $V^*(s)$, which is computed recursively using the Bellman equation:

$$V^*(s) = \max_{a \in A} \left[R(s, a) + \gamma \sum_{s'} T(s, a, s') V^*(s') \right], \tag{1}$$

where $\gamma \in [0, 1]$ is the discount factor. The optimal policy π^* is recoverable from $V^*(s)$.

B. Model-Based Prognostics Architecture

Given a system model defined as:

$$x(k+1) = f(k, x(k), \theta(k), u(k), v(k)), \tag{2}$$

$$y(k) = h(k, x(k), \theta, u(k), n(k)), \tag{3}$$

where k is the discrete time variable, $x(k) \in \mathbb{R}^{n_x}$ is the sate vector, $\theta(k) \in \mathbb{R}^{n_\theta}$ is unknown parameter vector, $u(k) \in \mathbb{R}^{n_u}$ is the input vector, $v(k) \in \mathbb{R}^{n_v}$ is the process noise vector, f is the state equation, $y(k) \in \mathbb{R}^{n_y}$ is the output vector, $n(k) \in \mathbb{R}^{n_n}$ is the measurement noise vector, and h is the output model.

For the system described above, a model-based prognostics architecture, adopted from [33], consists of two major steps: estimation and prediction. In the estimation step, the joint state-parameter estimate $p(x(k), \theta(k)|y(k_0 : k))$ is computed using the system dynamics and observation history up to time k, denoted as $y(k_0 : k)$. This step typically involves the use of a state estimator. In the prediction step, the probability distribution $p(k_E(k_P)|y(k_0 : k_P))$ at the prediction time k_p is computed using the joint state-parameter estimate along with hypothesized future inputs of the system.

IV. Methodology

In this section, we provide a high-level description of the problem formulation and the proposed framework, along with a discussion of the various interconnected components of the framework.

A. Problem Formulation

We consider a drone package delivery operation in unstructured airspace, where a homogeneous fleet of drones operates within a city-scale domain. The operation involves product depots with known geographic locations. The drones transport packages from these depots to delivery locations. Based on the payload capacities of drones [34], we assume that each drone carries one package at a time and is assigned to fly to a single destination.

The problem we address is ensuring the safe and efficient operation of these drones in terms of both collision avoidance and available battery energy. Therefore, a feasible solution to our problem must provide safe and energy-efficient trajectories for each aircraft and assess the mission's feasibility concerning battery capacity before departure.

B. Approach Overview

The overall objective of this paper is to introduce a trajectory planning and decision-making framework that facilitates the generation of energy-efficient and strategically de-conflicted flight plans, along with the assessment of their feasibility concerning potential battery energy insufficiency. To address this challenge, we present a comprehensive algorithmic framework, illustrated in Figure 2, which employs a two-layer approach.

The upper layer (Section VII) handles the trajectory planning task by generating safe and energy-efficient trajectories

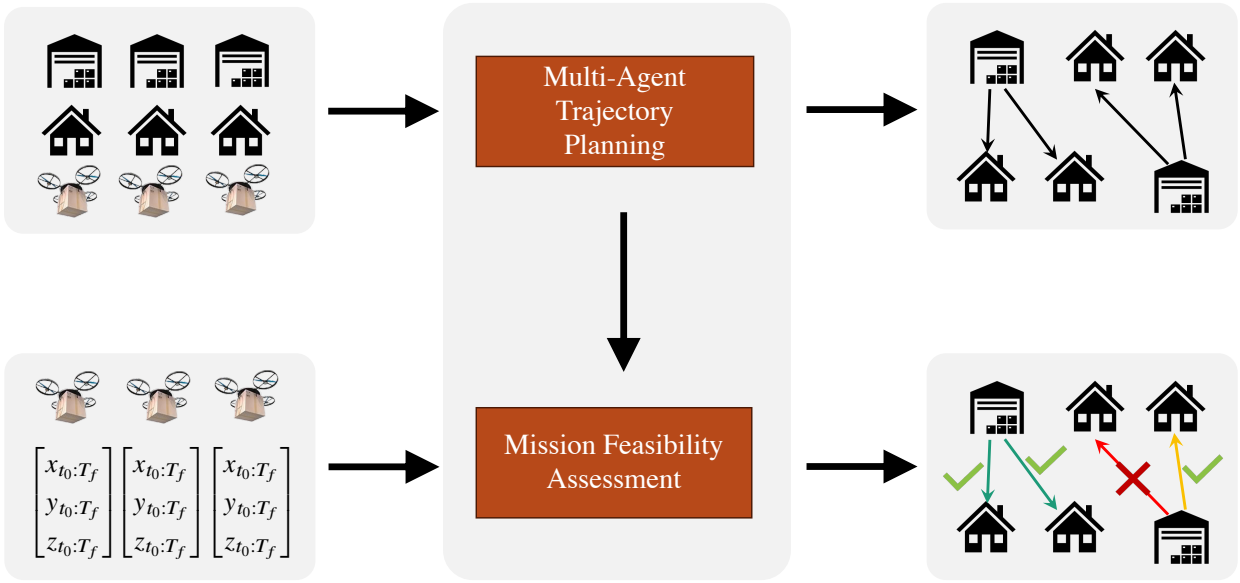

Fig. 2 We developed a two-layer framework to address the challenges of trajectory planning and battery energy-related feasibility assessment.

for each aircraft in the environment. It takes the depot-to-package delivery location assignments for each aircraft as input and employs an MDP-based multi-agent trajectory planning framework to generate trajectories that: (*i*) avoid collisions with other aircraft in the environment, (*ii*) are smooth and short to minimize excessive energy consumption, and (*iii*) are energy-efficient with respect to the wind field in the urban environment. This formulation inherently supports decentralized trajectory planning, as each aircraft optimizes its trajectory based on its own local reward function while still ensuring separation from others.

The lower layer (Section VII), focuses on feasibility assessment and decision-making regarding the execution of the mission. It takes the safe and energy-efficient trajectories generated by the upper layer and utilizes a model-based prognostic architecture to predict the final State of Charge (SoC) and quantify the uncertainty associated with each flight mission. Finally, the lower layer decides whether to proceed with the flight mission using the generated trajectory and at the assigned flight time.

V. Upper Layer: Trajectory Planner

The trajectory planner implemented in this paper, introduced in [35] and [20], represents the decision-making problem as a Markov decision process (MDP). A significant update to the current planner, as compared to [20], is the inclusion of energy considerations in the trajectory planning problem. Therefore, this paper focuses on highlighting the main components of the planner, with a particular emphasis on the energy-related reward function. It is also worth noting that, although the trajectory planner is designed for online guidance of aircraft in a free-flight manner, in this study, it has been employed offline to perform pre-departure flight planning.

A. Aircraft Dynamics

The state transition function $T(s_t, a, s_{t+1})$ of the MDP-based decision maker is given by a six-degree-of-freedom (6-DOF) low-fidelity dynamic model of octo-rotors, adopted from [36], where inertia is described by lumped masses. The model employs lumped masses to describe the inertia and provides a balance between accuracy and computational efficiency. As such, it accurately represents the behavior of the aircraft without imposing excessive computational load. The model is formulated in the state-space representation as shown in Equation 4, where s, c, and t denote sin, cos, and tan functions, respectively.

$$\dot{x} = \begin{bmatrix} \dot{x} \\ \dot{y} \\ \dot{z} \\ (s_\theta c_\psi c_\phi + s_\phi s_\psi)\frac{T}{m_t} \\ (s_\theta s_\psi c_\phi - s_\phi s_\psi)\frac{T}{m_t} \\ -g + c_\phi c_\theta \frac{T}{m_t} \\ p + q s_\phi t_\theta + r c_\phi t_\theta \\ q c_\phi - r s_\phi \\ q\frac{s_\phi}{c\theta} + r\frac{c_\phi}{c\theta} \\ \frac{I_{ybyb} - I_{zbzb}}{I_{xbxb}} qr + \frac{l}{I_{xbxb}}\tau_\phi \\ \frac{I_{zbzb} - I_{xbxb}}{I_{ybyb}} pr + \frac{l}{I_{ybyb}}\tau_\theta \\ \frac{I_{xbxb} - I_{ybyb}}{I_{zbzb}} pr + \frac{l}{I_{zbzb}}\tau_\psi \end{bmatrix} \tag{4}$$

In the above equation, linear position and velocity variables are represented by $s, \dot{s}$, and angular position and velocity variables by $[\phi, \theta, \psi]$ and $[p, q, r]$, respectively. The state vector is defined as $x = [s, \dot{s}, \phi, \theta, \psi, p, q, r]^T$. Additionally, g denotes the gravitational constant, m_t represents the lumped mass, and l is the arm length.

B. State Space

The environment is a continuous three-dimensional state space located in any chosen geographic region. Given the dynamics of an aircraft:

$$\dot{\zeta}(t) = f(\zeta(t), u(t)), \tag{5}$$

where, $f : \mathbb{R}^n \times \mathbb{R} \rightarrow \mathbb{R}^n$ is a continuous function. ζ denotes the aircraft states. The trajectory of an aircraft $\xi : \mathbb{R}^n \times \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}^n$ is the solution to the differential equation (5). For a given initial set $x_0 \in \mathbb{R}^n$, the state of the system at time t is $\xi(\zeta_0, t) = \zeta(t)$. The control input $u(t) \in \mathbb{R}^m$ is comprised of the thrust T, and the three torques τ_θ, τ_ψ, and τ_ψ. In addition, a single state in the state space (s_o) contains all the states of an aircraft (ζ) and the states of every other aircraft denoted as $f_j, \forall j \in J$. Thus, we can define s_o as $s_o = [\zeta, f_1, ..., f_j]$, where j represents the number of other aircraft.

C. Action Space

The action space of the MDP is composed of the individual action spaces of the four inputs: the thrust T, and three torques τ_θ, τ_ψ, and τ_ϕ. The action space of T is composed of 9 logarithmically spaced discrete values between $50N$ and $500N$. The chosen minimum and maximum values are determined by the performance capabilities of the motors (KDE-4213XF motors) in terms of the delivered thrust and torque. Consequently, the inputs of τ_θ, τ_ψ, and τ_ϕ are linearly spaced within a range of 7 input values.

The linearly spaced input set in Nm is constructed as follows:

$$\tau_\theta = [-150, -100, -50, 0, 50, 100, 150] \tag{6}$$

$$\tau_\psi = [-150, -100, -50, 0, 50, 100, 150] \tag{7}$$

$$\tau_\phi = [-150, -100, -50, 0, 50, 100, 150]. \tag{8}$$

Finally, the joint action space becomes:

$$\mathcal{A} = \{T, \tau_\theta, \tau_\psi, \tau_\phi\}. \tag{9}$$

VI. Reward Function Design

The reward function is the primary mechanism to control the behavior of an MDP agent. It is defined as $\mathcal{R}(s_t, a_t, s_{t+1})$, representing the reward an agent collects when it transitions from state s_t to s_{t+1} after taking action a_t. In this work, we have utilized both positive and negative rewards, as presented in Table 1, to guide the aircraft to their destination while avoiding potential collision with nearby aircraft. A detailed description of the methods used to design the collision avoidance reward function can be found in [35] and [20]. However, since this paper focuses on the energy efficiency of trajectories, the remainder of this section discusses the methods used to design the energy-related reward functions.

The goal of the energy-related reward function is to guide the trajectory planning agent in generating energy-efficient trajectories. Incorporating energy consumption into trajectory planning is essential because an aircraft's energy usage

depends heavily on the type of reference trajectories generated. These trajectories play a direct role in determining the feasibility of the flight mission. In this section, we outline four approaches for incorporating energy efficiency into trajectory generation: energy cost function, trajectory length, deviation from direct paths, and wind-aware planning. A schematic diagram that shows the overall idea behind these reward functions is presented in Figure 3.

A. Energy Cost Function

A key component of the energy-efficiency reward function is the learning-based energy cost function, represented by $\Phi(\cdot)$. This function estimates the energy consumption associated with each future state of the aircraft. Let $\mathbf{s}_t \in \mathbb{R}^3$ denote the aircraft's current position, and let $\mathbf{s}_{t+k,a} \in \mathbb{R}^3$ represent each future state of the aircraft across multiple time steps and actions, where k corresponds to the number of steps into the future from the current state $\mathbf{s}_t$, and a represents control actions. At each decision-making step, the energy cost function $\Phi(\cdot)$ takes the changes in aircraft positions ($\Delta\mathbf{d}_{t+k}$) as input to predict the energy required for the aircraft to transition from its current state to each projected future state. By using position changes as input, the function avoids dependency on the initial aircraft states, enabling efficient energy estimation. Mathematically, the energy consumption of the aircraft at future state $\mathbf{s}_{t+k,a}$ denoted $\mathbb{E}_{t+k}$ can be represented as:

$$\mathbb{E}_{t+k} = \Phi(\Delta\mathbf{d}_{t+k}; \theta) \tag{10}$$

where, $\Phi(\cdot)$ is the energy cost function, θ represents the parameters of the model, and the vector $\Delta\mathbf{d}_{t+k} = [\Delta d_1, \Delta d_2, \ldots, \Delta d_k]$ represents the change in distance traveled by the aircraft from its position at time t to the k^{th} future state and can be mathematically described as:

$$\Delta\mathbf{d}_{t+k} = \sum_{i=1}^{k} \sqrt{(x_i - x_{i-1})^2 + (y_i - y_{i-1})^2 + (z_i - z_{i-1})^2}. \tag{11}$$

Given a dataset $\mathcal{D}$ consisting of N samples $\{(\Delta\mathbf{d}_{t+k}^{(i)}, \mathbb{E}_{t+k}^{(i)})\}_{i=1}^{N}$, where each sample includes a change in distance and the corresponding observed energy consumption, the learning task is to minimize the prediction error of the model. This can be formalized as:

$$\theta^* = \arg\min_{\theta} \sum_{i=1}^{N} \left(\mathbb{E}_{t+k}^{(i)} - \Phi(\Delta\mathbf{d}_{t+k}^{(i)}; \theta)\right)^2. \tag{12}$$

To this end, we trained a regression model given in Equation 13 that predicts the energy cost ($\mathbb{E}_{t+k}$) based on changes in aircraft positions.

$$\mathbb{E}_{t+k} = \mathbf{A}^\top \cdot \Delta\mathbf{d}_{t+k} + \mathbf{B}^\top, \tag{13}$$

where, $\mathbf{A}^\top$ and $\mathbf{B}^\top$ are the transposed coefficient vectors of the energy cost function, they are defined as follows:

$$\mathbf{A} = [0, 520.699, 203.742, 127.606, 86.3024, 71.8159, 63.1053, 57.9294, 50.7294, 46.399],$$

$$\mathbf{B} = [0, 369.168, 869.01, 1119.61, 1459.52, 1368.97, 1279.89, 1128.54, 1144.29, 1183.82].$$

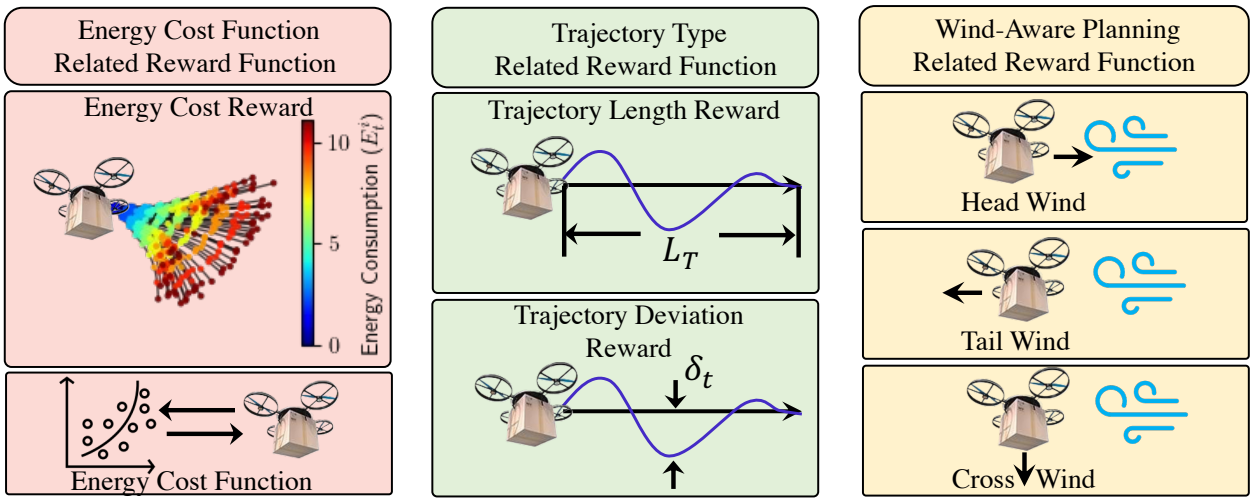

Fig. 3 The energy-efficiency reward function consists of three components: the energy cost reward, the trajectory type reward, and the wind-aware planning reward.

Finally, the energy-related reward ($\mathcal{R}_{\text{energy}}$) at time t can be expressed as:

$$\mathcal{R}_{\text{energy}} = -w_{\text{E}} \cdot (\mathbf{A}^\top \cdot \Delta \mathbf{d}_{t+k} + \mathbf{B}^\top). \tag{14}$$

The negative sign in the above equation indicates that higher energy consumption leads to a lower reward, thereby encouraging the aircraft to select actions that minimize energy usage. The term w_{E} is a constant scaling factor.

B. Trajectory Deviation Reward

The trajectory deviation reward is another component of the energy efficiency reward function, designed to minimize deviations from the straight-line trajectory that directly connects the aircraft's initial position and its goal. A straight-line path is generally the most energy-efficient route, as it minimizes unnecessary maneuvers that can increase energy consumption. By penalizing deviations from this ideal path, the reward function encourages the aircraft to maintain a smooth and efficient trajectory throughout its flight.

The overall idea of the trajectory deviation reward function is to compute the angle between a straight line that connects the goal state, denoted as $\mathbf{s}_G \in \mathbb{R}^3$, and current position, $\mathbf{s}_t \in \mathbb{R}^3$ and each future state of the aircraft $\mathbf{s}_{t+k,a} \in \mathbb{R}^3$. To that end, we first compute the goal vector $\mathbf{g}$, which is the direction from the current position $\mathbf{s}_t$ to the goal $\mathbf{s}_G$. This

vector is obtained as the difference between the two positions, $\mathbf{s}_G$ and $\mathbf{s}_t$. Then, to facilitate comparisons with other vectors, the goal vector is converted to a unit vector $\hat{\mathbf{g}}$. This procedure can be described mathematically as:

$$\hat{\mathbf{g}} = \frac{\mathbf{g}}{\|\mathbf{g}\|} = \frac{\mathbf{s}_G - \mathbf{s}_0}{\sqrt{g_x^2 + g_y^2 + g_z^2}}. \tag{15}$$

Similarly, for each possible future state, a unit vector that represents the direction from the current state to the future state $\hat{\mathbf{v}}_{t,a}$ is computed as:

$$\hat{\mathbf{v}}_{t+k,a} = \frac{\mathbf{v}_{t+k,a}}{\|\mathbf{v}_{t+k,a}\|} = \frac{\mathbf{s}_{t+k,a} - \mathbf{s}_t}{\sqrt{v_{t+k,a,x}^2 + v_{t+k,a,y}^2 + v_{t+k,a,z}^2}}. \tag{16}$$

Once these vectors are computed, the angle between the vectors leading to future states and the goal direction $\theta_{t,a}$ can be found using the dot product between each normalized future state vector $\hat{\mathbf{v}}_{t+k,a}$ and the normalized goal vector $\hat{\mathbf{g}}$. The result of this dot product gives the cosine of the angle $\theta_{t+k,a}$ between the vectors. Mathematically, this is expressed as:

$$\theta_{t+k,a} = \cos^{-1}(\hat{\mathbf{v}}_{t+k,a} \cdot \hat{\mathbf{g}}). \tag{17}$$

Finally, the trajectory deviation reward, $\mathcal{R}_{\text{deviation}}$, is then defined as the negative of the deviation at each time step multiplied with a scaling factor w_D:

$$\mathcal{R}_{\text{deviation}} = -w_D \cdot \theta_{t+k,a}. \tag{18}$$

This formulation ensures that the aircraft is encouraged to follow the most direct and energy-efficient path from its initial position to its goal, minimizing unnecessary deviations that could lead to increased energy consumption.

C. Trajectory Length Reward

The trajectory length reward is designed to minimize the total distance traveled by the aircraft from the initial position to the goal. The key difference between this and the trajectory deviation reward is that the latter focuses on maintaining alignment with the straight-line reference path — a distinction that is particularly significant in the presence of dynamic obstacles or wind disturbances. To accurately account for the entire trajectory, the path is divided into three distinct segments: the past trajectory (L_{past}), the current decision-making horizon (L_{horizon}), and the remaining distance to the goal (L_{goal}). This segmentation enables a comprehensive evaluation of the entire trajectory, thereby encouraging the aircraft to minimize the length of its path across the full mission duration. The total path length L_{total} for the trajectory from the initial position $\mathbf{s}_0$ to the goal $\mathbf{s}_G$is, therefore, the sum of these three segments. This can be expressed mathematically as:

$$L_{\text{total}} = L_{\text{past}} + L_{\text{horizon}} + L_{\text{goal}}, \tag{19}$$

where, L_{total} represents the distance traveled by the aircraft from the initial position $\mathbf{s}_0$ to the current state $\mathbf{s}_t$, L_{horizon} represent the predicted distance within the current decision-making horizon, extending from the current state $\mathbf{s}_t$ to the state at K^{th} future step, and L_{goal} represent the distance from the predicted state at time k to the goal $\mathbf{s}_G$. These are defined mathematically as:

$$L_{\text{past}} = \sum_{i=1}^{t} \|\mathbf{s}_i - \mathbf{s}_{i-1}\|, \; L_{\text{horizon}} = \sum_{k=t+1}^{K} \|\mathbf{s}_{t+k} - \mathbf{s}_{t+k-1}\|, \text{ and } L_{\text{goal}} = \|\mathbf{s}_G - \mathbf{s}_K\|. \tag{20}$$

Accordingly, the trajectory length reward R_{Length} is defined as the negative of the total trajectory length:

$$\mathcal{R}_{\text{Length}} = -w_{\text{L}} \cdot L_{\text{total}}, \tag{21}$$

where w_{L} is a constant scaling factor. This formulation effectively captures the entire trajectory from the initial position to the goal and incentivizes the aircraft to minimize the overall path length and enhance energy efficiency throughout the mission.

D. Wind-Aware Planning Reward

In designing energy-efficient trajectories, considering wind effects on energy consumption is essential. The wind-aware planning reward function aims to generate energy-efficient paths by encouraging the aircraft to seek tailwinds and avoid headwinds. The main idea is to minimize energy usage by aligning the aircraft's velocity with favorable wind directions.

A key insight behind the formulation of this reward function is the angle χ between the wind vector and the aircraft's velocity vector, which significantly impacts energy requirements. When χ is close to 0 radians (headwind scenario), the wind opposes the aircraft's travel direction, increasing air resistance and energy consumption. Conversely, when χ is near π radians (tailwind scenario), the wind aligns with the aircraft's direction, reducing energy consumption. In crosswind conditions, where $\chi = \frac{\pi}{2}$ or $\chi = \frac{3\pi}{2}$, the wind blows perpendicular to the aircraft's path, resulting in moderate energy consumption.

Let $\mathbf{v}_{\text{aircraft}} = [v_x, v_y]$ represent the velocity vector of the UAV, where v_x and v_y are the components of the UAV's speed in the x and y directions, respectively. Let $\mathbf{w} = [w_x, w_y]$ denote the wind speed vector, where w_x and w_y are the wind speed components in the x and y directions. Finally, let $\mathbf{p} = [x, y]$ denote the position of the UAV in the horizontal plane. The wind-aware reward function is computed at each position $\mathbf{p}$, based on the relative orientation between the UAV's velocity vector and the wind vector at that location.

To compute the angle between the UAV's velocity vector $\mathbf{v}_{\text{aircraft}}$ and the wind vector $\mathbf{w}$, we use the dot product

between the two vectors:

$$\cos(\chi) = \frac{\mathbf{v}_{\text{aircraft}} \cdot \mathbf{w}}{\|\mathbf{v}_{\text{aircraft}}\|\|\mathbf{w}\|}, \tag{22}$$

where $\mathbf{v}_{\text{aircraft}} \cdot \mathbf{w} = v_x w_x + v_y w_y$, $\|\mathbf{v}_{\text{aircraft}}\| = \sqrt{v_x^2 + v_y^2}$ and $\|\mathbf{w}\| = \sqrt{w_x^2 + w_y^2}$ are the norms of the UAV's velocity and wind vectors, respectively. The above dot product quantifies how well the UAV's velocity aligns with the wind direction. A positive value indicates a tailwind effect, while a negative value indicates a headwind effect. Finally, the wind-aware planning reward is computed as:

$$\mathcal{R}_{\text{wind}} = -w_{\text{W}} \left(\mathbb{B} - \mathbb{A}\cos(\chi)\right), \tag{23}$$

where w_{W} is a scaling factor, and $\mathbb{A}$ and $\mathbb{B}$ are constants that modulate the reward. The constant $\mathbb{A}$ controls the amplitude of the reward function, with larger values of $\mathbb{A}$ enhancing the difference in penalty between headwind and tailwind scenarios, thus penalizing headwinds more severely while reducing the penalty for tailwinds. The constant $\mathbb{B}$ adjusts the baseline of the reward function, ensuring that crosswinds ($\chi \approx \frac{\pi}{2}$) are also penalized to an appropriate degree.

E. Combining the Reward Functions

Now that we have discussed the formulation of each energy-efficiency-related reward function, we can combine these functions to enable our trajectory planning agent to generate safe and energy-efficient trajectories, expressed as:

$$\mathcal{R}_{\text{total}} = \mathcal{R}_{\text{destination}} + \mathcal{R}_{\text{separation}} + \mathcal{R}_{\text{efficiency}}, \tag{24}$$

where $\mathcal{R}_{\text{destination}}$ and $\mathcal{R}_{\text{separation}}$ represent rewards that guide the aircraft to its assigned destination while maintaining safe separation from other aircraft. These reward functions are discussed in [20] and [35], with details presented in Table 1. The separation is ensured using reachability analysis, which computes the set of states reachable by each aircraft over a given time horizon and identifies potential conflicts. By penalizing states that overlap with these reachable sets, the framework guarantees collision avoidance while enabling efficient trajectory planning. The term $\mathcal{R}_{\text{efficiency}}$ denotes the sum of energy-efficiency-related reward functions discussed in this section, defined as:

$$\mathcal{R}_{\text{efficiency}} = w_{\text{L}}\mathcal{R}_{\text{length}} + w_{\text{D}}\mathcal{R}_{\text{deviation}} + w_{\text{E}}\mathcal{R}_{\text{energy}} + w_{\text{W}}\mathcal{R}_{\text{wind}}. \tag{25}$$

The weights w_{L}, w_{D}, w_{E}, and w_{W} are tuning parameters that control the relative importance of each reward function in the overall trajectory planning process. In this work, we set these weights as $w_{\text{L}} = 0.1$, $w_{\text{D}} = 0.5$, $w_{\text{E}} = 0.2$, and $w_{\text{W}} = 0.2$, reflecting a higher emphasis on minimizing trajectory deviation while balancing trajectory length, energy consumption, and wind influence.

Table 1 Reward Functions for Each Aircraft

Reward Source	Reward Function	Location	Decay Factor	Purpose
Intruder aircraft	$-(100t + 500)$	Intruder	0.97	Collision avoidance
Destination	200	Manually placed	0.999	Destination attraction
Ownship	$-w_{\mathrm{E}} \cdot (\mathbf{A}^{\top} \cdot \Delta \mathbf{d}_{t+k} + \mathbf{B}^{\top})$	Aircraft next states	0	Energy efficiency
Ownship	$-w_{\mathrm{D}} \cdot \theta_{t+k,a}$	Aircraft next states	0	Energy efficiency
Ownship	$-w_{\mathrm{L}} \cdot L_{\mathrm{total}}$	Aircraft next states	0	Energy efficiency
Ownship	$-w_{\mathrm{W}} (\mathbb{B} - \mathbb{A} \cos(\chi))$	Aircraft next states	0	Energy efficiency

VII. Lower Layer: Feasibility Assessment and Decision Making

A key task in assessing the feasibility of a mission related to battery energy is predicting the battery's state of charge (SoC) at the end of the flight mission. The SoC of a battery is typically defined as 1 when fully charged and 0 when discharged to a predetermined voltage threshold. By obtaining the ending SoC with certain probabilistic guarantees, we can evaluate the mission's feasibility by comparing this prediction against system requirements and deciding whether to execute the mission. To achieve this, we adopted a model-based prognostics architecture from [33] and its implementation from [37], whose basic working principles are discussed in Section III.B.

The battery model utilized in this study is an electrochemical-based model of Lithium-ion batteries, as described in [33], which are a popular choice for powering unmanned aerial vehicles. In this model, the battery's current draw (I) serves as the input, while the battery temperature (tb) and the voltage drop caused by solid-phase ohmic resistance (V_o) represent its outputs ($y(k)$). The estimation algorithm used in this paper is the Unscented Kalman Filter (UKF) [38], along with the battery model.

The predictor algorithm used in this paper is the Monte Carlo predictor [39], which randomly samples from the battery's current state distribution, and each sample is simulated to the end of the flight. By collecting a set of SoC values from several Monte Carlo simulations, the probability distribution can be built, and the probability of mission success at a given time t can be computed using the following equation:

$$P_{\mathrm{success}}(t) = \frac{\sum_{i=1}^{n} (\mathrm{SoC}_i(t) > \mathrm{SoC}_{\mathrm{th}})}{n}, \tag{26}$$

where n represents the number of Monte Carlo simulations in the prediction step and $\mathrm{SoC}_{\mathrm{th}}$ is the threshold battery SoC value required at the end of the flight.

The battery SoC prediction process is illustrated in Figure 4. According to the procedure, given a reference trajectory and information about the available battery energy, the first step is to simulate the high-fidelity octo-rotor model adopted from [40] along the trajectory to obtain the mission's current requirement (I_{req}), which is used as a future load for battery SoC prediction. Once the mission's I_{req} is known, the battery model will be simulated from its initial state until SoC is reached. Finally, the predicted SoC points are collected, and the probability density function for battery SoC at

the flight end time is constructed.

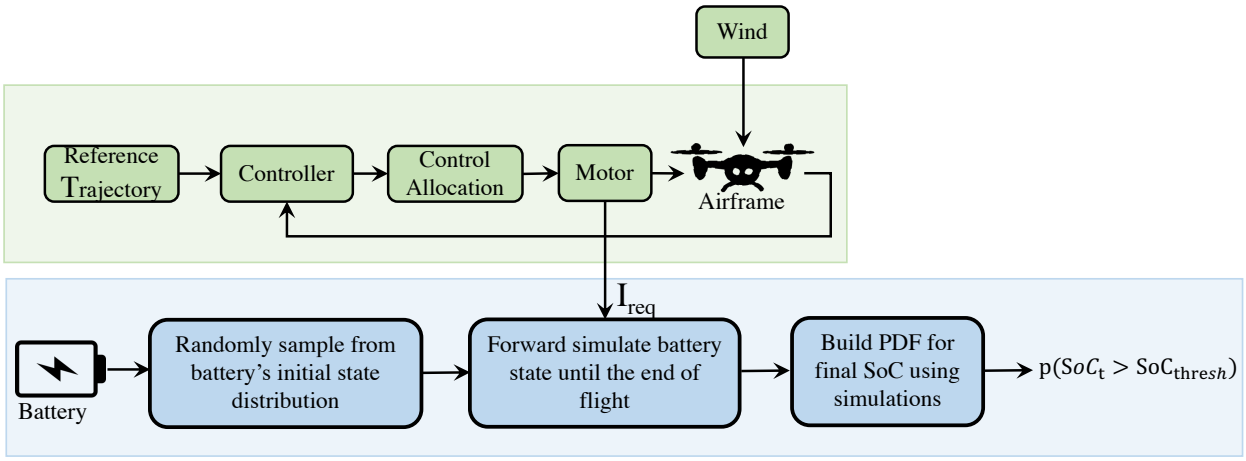

Fig. 4 Schematic diagram representation of the feasibility assessment procedure.

VIII. Overall Framework

The overall working procedure of the framework is presented in Algorithm 1. The inputs to the framework are the aircraft's initial state and destination, the battery's initial state (x_0), and mission success probability threshold $P_{\text{threshold}}$ to determine whether to conduct the mission or not. The framework's output is the decision to either take off or hold the aircraft. The first step is generating trajectories from start to destination for each aircraft. These trajectories are safe (in terms of collision avoidance with other aircraft) and locally energy optimal. Next, before proceeding to the battery prognostics process, we sample waypoints from the generated trajectories and feed those waypoints to the detailed aircraft dynamics as a reference trajectory. Then, we run the detailed octo-rotor dynamics and collect the current requirements of the missions. Once the current requirement is collected, we randomly sample from the battery's initial state and simulate the battery model until the end of the flight is reached.

IX. Experiments and Results

A. Package Delivery Scenario Description

We designed a drone package delivery scenario in a region near the city of Boston around coordinates 42°15′17.9″N and 71°08′27.7″W, to demonstrate the performance of the proposed framework. Figure 5 shows the Google Earth view of the area. Since the drones are expected to operate within an urban setting where buildings and other static structures influence airflow patterns, it is essential to use CFD-simulated wind data that captures localized variations in wind speed and direction. To achieve this, we incorporated CFD-simulated wind data for the area from [41], as shown in Figure 6. The wind data shown in Figure 6 is a result of simulation for a wind speed of approximately $10.2 m/s$, heading

Algorithm 1: Pre-departure Flight Planning and Mission Feasibility Assessment Framework

Procedure `FlightPlanAndRiskAssessment()`:
- **Input :** Aircraft initial state and destination, Battery initial state (x_0) and $P_{\text{threshold}}$
- **Output :** Decision to either take off or hold the aircraft
1. **for** *each aircraft i* **do**
2. Generate safe and energy efficient trajectory from start to destination
3. Obtain the current requirement of the trajectory
4. Perform Monte Carlo (MC) simulations of the battery using the current requirement
5. Construct End-of-Flight Battery SoC Probability Density Function (PDF) from MC simulation results
6. Compute the probability of mission success using Equation 26
7. **if** $P_{success} > P_{threshold}$ **then**
   - Execute the mission
8. **else**
   - Keep charging the aircraft and/or re-plan the flight trajectory
9. End of operation

at zero degrees from left to right across the domain.

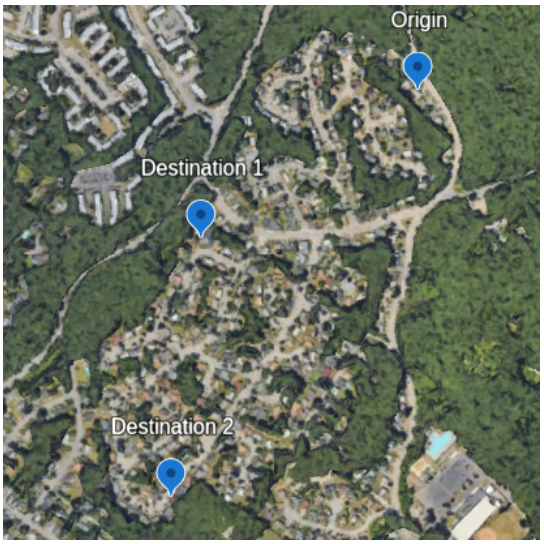

Fig. 5 Google Earth View of the Area.

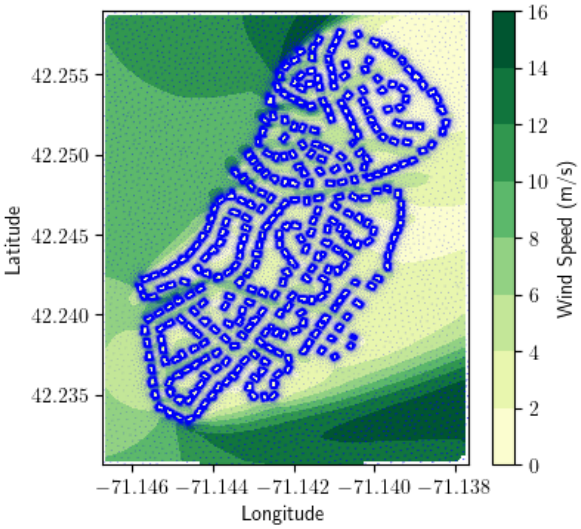

Fig. 6 CFD-Simulated Wind Field for the Area.

We then designate a depot location, which serves as a hub where drones are stationed for charging, maintenance, and package loading. This depot is the operational center from which drones depart to deliver packages and return after completing their missions. In our scenario, the depot is assumed to be located at 42°15′21.4″N and 71°08′20.1″W. From this depot, multiple package delivery missions can be assigned to destinations within the area. A mission is considered successful when the assigned aircraft takes off with the package, flies to the destination, descends to the drop-off point, releases the package, and returns safely to the depot. Our objective is to use the proposed framework to (1) generate safe and energy-efficient trajectories and (2) assess the feasibility of completing these missions with the available energy.

B. Simulation Results

Although the scenario can accommodate a configurable number of aircraft that use the proposed planning and decision-making framework, we present experimental results for a specific case involving two aircraft. These aircraft are assigned to take off from the depot and fly to two separate destinations located at 42°15′13.0′′N and 71°08′36.6′′W and 42°14′58.5′′N and 71°08′38.9′′W, respectively. As such, the two flight trajectories generated by the upper layer of the framework are shown in Figures 7a and 7b. In each figure, two sets of trajectories illustrate the impact of the energy-efficient reward function. The red dashed lines represent trajectories generated without considering energy efficiency, while the blue trajectories account for energy-efficient planning. As shown, the energy-efficient trajectories are smoother, more direct, and maintain a consistent altitude during cruising. Consequently, as illustrated in Figures 7c and 7d, incorporating energy-efficient rewards significantly reduces both the power consumption and cumulative energy consumption for both missions.

Once these trajectories are generated, we proceed to the lower layer of the proposed framework to perform prognostics and mission feasibility assessment. We generated 500 Monte Carlo simulations to assess the feasibility of each mission. Additionally, to determine the probability of mission success or failure, we established a threshold State of Charge (SoC_{th}) value of 50% at the end of the flight. Figure 8 shows the SoC and voltage predictions for the two flight missions mentioned above. From the two voltage plots in Figures 8d and 8d, we observe a noticeable change in the voltage profiles in the middle of the simulations. This change is primarily due to the aircraft transitioning between flight phases—from cruising to descending to deliver the package, and then ascending again to return to its depot. Finally, to decide whether to conduct the mission or not, we used Equation 26 along with the constructed SoC PDF, SoC_{th}, and a probability threshold $P_{\text{threshold}}$ value of 95%. As mentioned in Algorithm 1, this $P_{\text{threshold}}$ value is used to determine the execution of the mission. Specifically, in our case, if the probability of an aircraft arriving at its depot after completing the delivery with a SoC value of 50% is above 95%, then our framework will decide to execute the mission. Accordingly, the final decisions for the mission shown in Figures 7a and 7b are given in Table 2.

Table 2 Mission Success Probabilities and Decisions

Mission ID	Success Probability	Decision
1	1.00	Cleared for Flight
2	0.08	Hold

C. Battery Feasibility Assessment Verification

In addition to the simulation-based experiments discussed in the previous section, we also conducted a hardware flight test to study the performance and accuracy of the feasibility assessment framework. We used an Octocopter aircraft shown in Figure 9 to conduct the actual flight test which is equipped with onboard sensors to collect real-time data on the

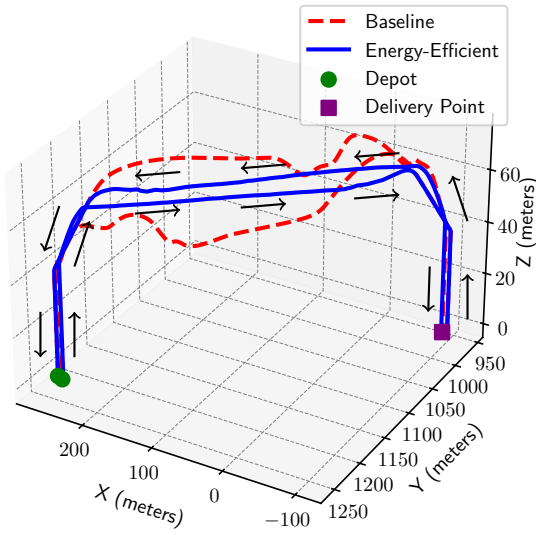

(a) Aircraft 1 Trajectory Comparison.

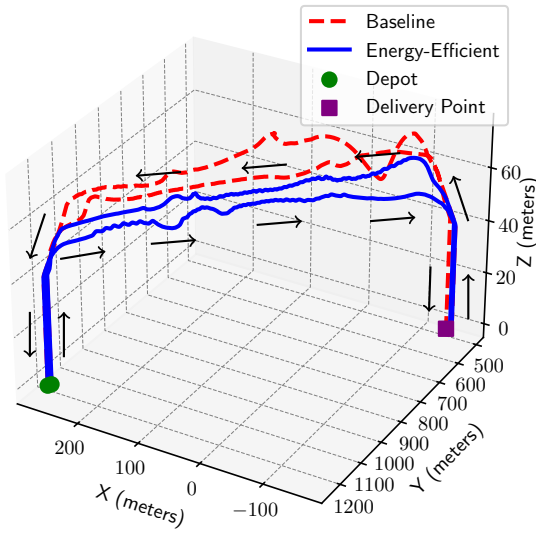

(b) Aircraft 2 Trajectory Comparison.

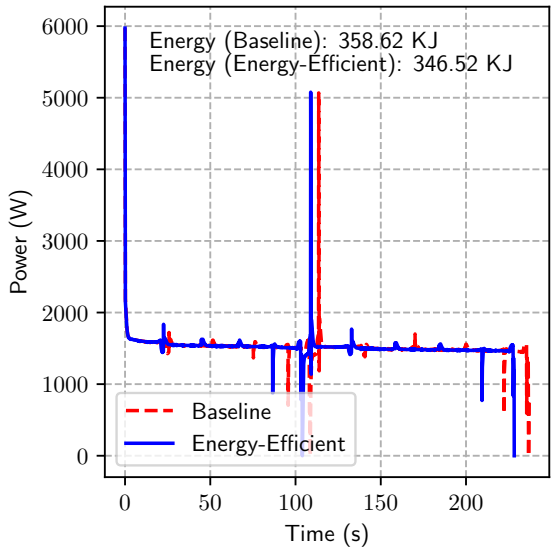

(c) Power Profile Comparison for Aircraft 1.

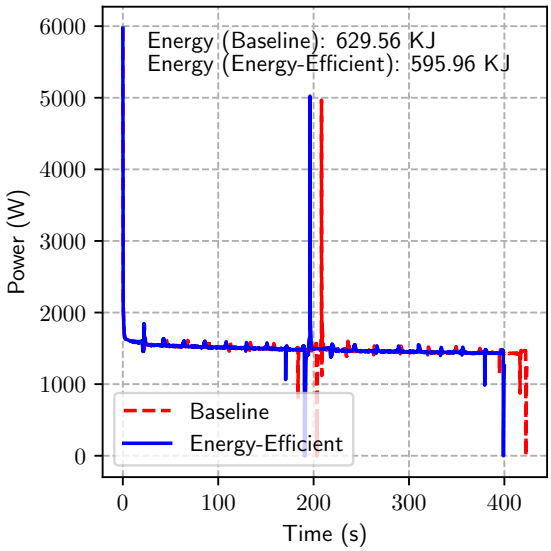

(d) Power Profile Comparison for Aircraft 2.

Fig. 7 Trajectory and Power Profile Comparison for Both Aircraft.

aircraft's battery-related parameters such as current, battery, and SoC to perform the validation. Figure 10 shows a flight plan at Poplar Ford Quinn Field test site with the first waypoint (waypoint1) indicating (38°50′6.9″N, 77°30′22.4″W) latitude and longitude coordinates. The maximum wind speed at the time of the flight experiment was close to 5 m/s.

After the flight test was conducted, we collected the current draw of the battery, shown in Figure 11. The figure illustrates power consumption variations across different flight phases, peaking during takeoff and spiking during aggressive maneuvers or wind disturbances, followed by a reduction during landing. The goal is to evaluate the

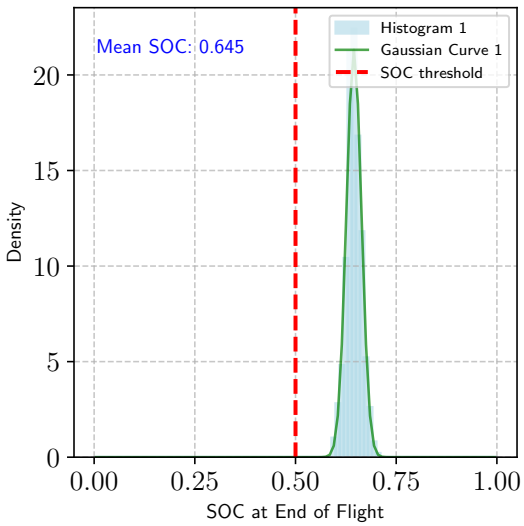

(a) SoC Prediction for Aircraft 1.

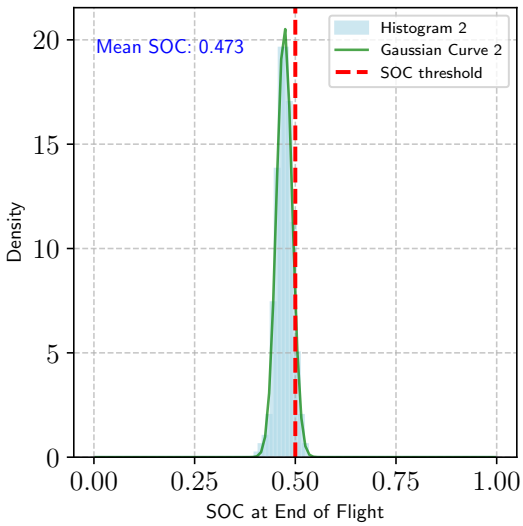

(b) SoC Prediction for Aircraft 2.

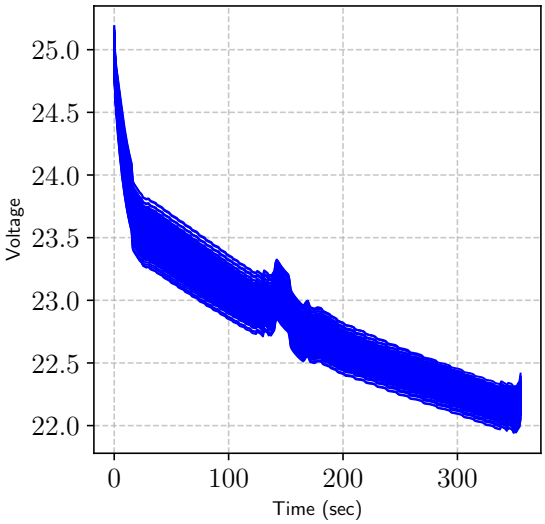

(c) Voltage Prediction for Aircraft 1.

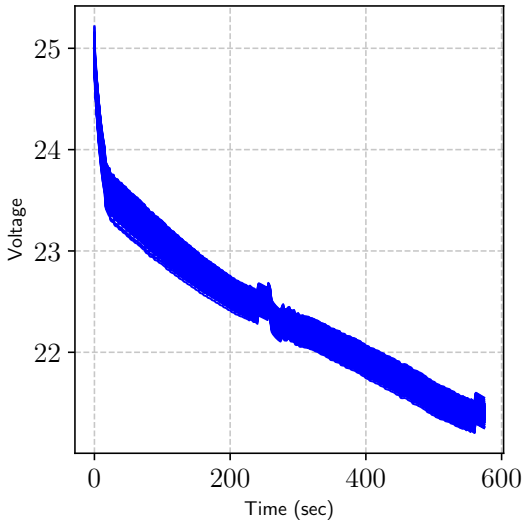

(d) Voltage Prediction for Aircraft 2.

Fig. 8 Battery Voltage and SOC Predictions for Both Aircraft.

prediction accuracy of the feasibility assessment framework using this data. To this end, we apply the collected current to the feasibility assessment framework and compare the predicted voltage and SoC trajectories against the actual voltage and SoC trajectories of the battery. However, since we are using a Monte Carlo-based feasibility assessment framework, we take the average of all voltage and SoC predictions to validate the performance of the framework, and the results are plotted in Figure 11.

As can be seen in the above Figure, the voltage and SoC trajectories we obtained from the feasibility assessment align well with the actual trajectories. To quantify the performance, we use both Root Mean Square Error (RMSE) and Mean Absolute Error (MAE) as metrics. RMSE measures the average magnitude of errors, giving more weight to larger discrepancies, thereby indicating how closely the simulated voltage and SoC trajectories match the actual flight data. In contrast, MAE provides a more balanced view of prediction errors, being less sensitive to outliers, which allows us to

Fig. 9 Snapshot of the Actual Flight.

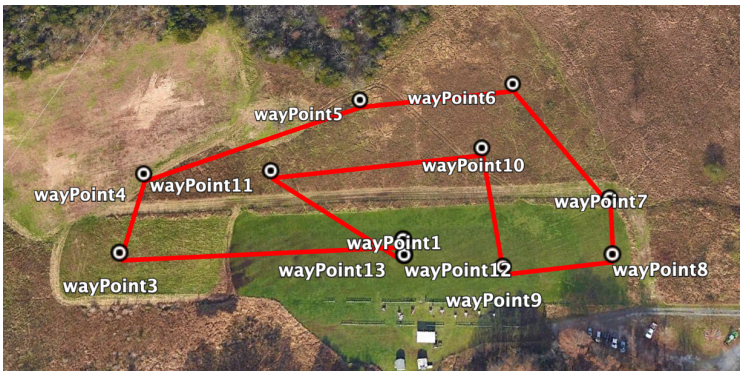

Fig. 10 Actual Flight Plan at Poplar Ford Quinn Field Test Site.

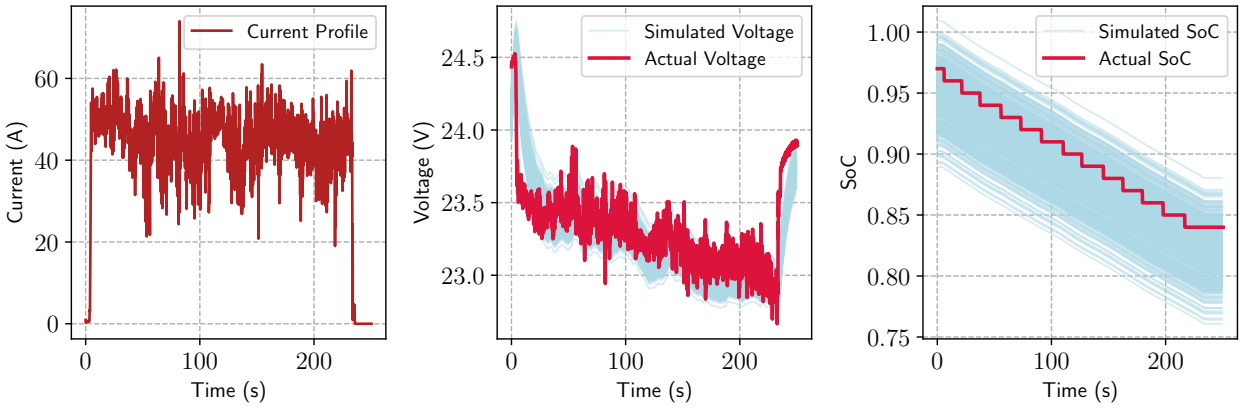

Fig. 11 Current profile of the mission, voltage comparison, and SoC comparison. The plots highlight the alignment between simulated results and actual measurements.

assess the typical deviation between the simulated and actual trajectories.

We calculated both RMSE and MAE separately for both voltage and SoC trajectories using the following equations:

$$\text{RMSE}_{\text{voltage}} = \sqrt{\frac{1}{N}\sum_{i=1}^{N}(V_i - \hat{V}_i)^2}, \quad \text{RMSE}_{\text{SoC}} = \sqrt{\frac{1}{N}\sum_{i=1}^{N}(\text{SoC}_i - \hat{\text{SoC}}_i)^2}, \tag{27}$$

where V_i and SoC_i represent the actual voltage and SoC values at time step i, and $\hat{V}_i$ and $\hat{\text{SoC}}_i$ are the corresponding predicted values. The Mean Absolute Error (MAE) for the voltage and SoC trajectories is defined as:

$$\text{MAE}_{\text{voltage}} = \frac{1}{N}\sum_{i=1}^{N}|V_i - \hat{V}_i|, \quad \text{MAE}_{\text{SoC}} = \frac{1}{N}\sum_{i=1}^{N}|\text{SoC}_i - \hat{\text{SoC}}_i|. \tag{28}$$

Finally, the $\text{RMSE}_{\text{voltage}}$ and RMSE_{SoC} values are 0.3377 and 0.1731, respectively, and the $\text{MAE}_{\text{voltage}}$ and MAE_{SoC} values are 0.0171 and 0.0167 for the experiment we conducted. This indicates that the feasibility assessment scheme achieves a high level of accuracy, with only minor deviations between the simulated and actual trajectories. The relatively

low RMSE values suggest that large errors are rare, while the lower MAE values reflect consistently small discrepancies, demonstrating the framework's ability to provide reliable predictions. These results highlight the effectiveness of the proposed framework in conducting accurate feasibility assessments, which are crucial for practical drone package delivery operations in real-world scenarios.

X. Conclusion

In conclusion, this paper presents a framework aimed at generating energy-efficient trajectory planning and feasibility assessment for drone package deliveries in urban environments. Developed to integrate with the UTM architecture, the framework effectively tackles the computational complexities of the problem by combining an MDP-based trajectory planner with a battery state of charge (SoC) prediction-based feasibility assessment in a two-layered approach. The upper layer generates energy-efficient flight paths by accounting for factors such as aircraft energy consumption model, trajectory types, and wind interactions, while the lower layer evaluates these trajectories for energy feasibility using an SoC-based uncertainty quantification approach.

We demonstrate the efficacy of the proposed framework in a realistic drone package delivery scenario set in an urban environment in Boston City. In addition, an actual flight experiment was conducted to validate the framework's performance. The results highlight the framework's capability to generate energy-efficient trajectories and assess the feasibility of each mission. Future works will explore the integration of learning-based online feasibility assessment and contingency trajectory planning to enhance decision-making capabilities during UAS operations. Specifically, contingency planning involves dynamically generating alternative trajectories in response to unexpected events, such as wind disturbances, obstacles, or energy constraints, ensuring that the aircraft can safely and efficiently complete its mission despite uncertainties.

Acknowledgments

This project is supported by NASA Grant 80NSSC21M0087 under the NASA System-Wide Safety (SWS) program.

References

[1] Federal Aviation Administration, "Package Delivery by Drone (Part 135)," , 2023. URL https://www.faa.gov/uas/advanced_operations/package_delivery_drone, accessed: 2024-10-19.

[2] Federal Aviation Administration, "FAA Makes Drone History in Dallas Area," , 2024. URL https://www.faa.gov/newsroom/faa-makes-drone-history-dallas-area, accessed: 2024-10-19.

[3] Marquand, Y. L., "FAA authorises Zipline and Wing for BVLOS operations in Dallas," , 2023. URL https://www.revolution.aero/news/2024/07/30/faa-authorises-zipline-and-wing-for-bvlos-operations-in-dallas/, accessed: 2024-10-19.

[4] Federal Aviation Administration, "Unmanned Aircraft System (UAS) Traffic Management (UTM) Concept of Operations," Tech. rep., U.S. Department of Transportation, 2020. URL https://www.faa.gov/sites/faa.gov/files/2022-08/UTM_ConOps_v2.pdf, version 2.0.

[5] Frazzoli, E., Mao, Z.-H., Oh, J.-H., and Feron, E., "Resolution of conflicts involving many aircraft via semidefinite programming," *Journal of Guidance, Control, and Dynamics*, Vol. 24, No. 1, 2001, pp. 79–86. https://doi.org/10.2514/2.4678.

[6] Raghunathan, A. U., Gopal, V., Subramanian, D., Biegler, L. T., and Samad, T., "Dynamic optimization strategies for three-dimensional conflict resolution of multiple aircraft," *Journal of guidance, control, and dynamics*, Vol. 27, No. 4, 2004, pp. 586–594. https://doi.org/10.2514/1.11168.

[7] Enright, P. J., and Conway, B. A., "Discrete approximations to optimal trajectories using direct transcription and nonlinear programming," *Journal of Guidance, Control, and Dynamics*, Vol. 15, No. 4, 1992, pp. 994–1002. https://doi.org/10.2514/3.20934.

[8] Schouwenaars, T., De Moor, B., Feron, E., and How, J., "Mixed integer programming for multi-vehicle path planning," *Proceedings of the 2001 European Control Conference (ECC)*, IEEE, 2001, pp. 2603–2608. https://doi.org/10.23919/ECC.2001.7076321.

[9] Richards, A., and How, J. P., "Aircraft trajectory planning with collision avoidance using mixed integer linear programming," *Proceedings of the 2002 American Control Conference (IEEE Cat. No. CH37301)*, Vol. 3, IEEE, 2002, pp. 1936–1941. https://doi.org/10.1109/ACC.2002.1023918.

[10] Pallottino, L., Feron, E. M., and Bicchi, A., "Conflict resolution problems for air traffic management systems solved with mixed integer programming," *IEEE Transactions on Intelligent Transportation Systems*, Vol. 3, No. 1, 2002, pp. 3–11. https://doi.org/10.1109/6979.994791.

[11] Vela, A., Solak, S., Singhose, W., and Clarke, J.-P., "A mixed integer program for flight-level assignment and speed control for conflict resolution," *Proceedings of the 48th IEEE Conference on Decision and Control (CDC) held jointly with 2009 28th Chinese Control Conference*, IEEE, 2009, pp. 5219–5226. https://doi.org/10.1109/CDC.2009.5400520.

[12] Delahaye, D., Peyronne, C., Mongeau, M., and Puechmorel, S., "Aircraft conflict resolution by genetic algorithm and B-spline approximation," *Proceedings of ENRI Int. Workshop on ATM/CNS. Tokyo, Japan*, Vol. 4, 2010, pp. 71–77.

[13] Hong, D., Lee, S., Cho, Y. H., Baek, D., Kim, J., and Chang, N., "Energy-efficient online path planning of multiple drones using reinforcement learning," *IEEE Transactions on Vehicular Technology*, Vol. 70, No. 10, 2021, pp. 9725–9740. https://doi.org/10.1109/TVT.2021.3109980.

[14] Razzaghi, P., Tabrizian, A., Guo, W., Chen, S., Taye, A., Thompson, E., Bregeon, A., Baheri, A., and Wei, P., "A survey on reinforcement learning in aviation applications," *Engineering Applications of Artificial Intelligence*, Vol. 136, 2024, p. 108911. https://doi.org/10.1016/j.engappai.2024.108911.

[15] Pontani, M., and Conway, B. A., "Particle swarm optimization applied to space trajectories," *Journal of Guidance, Control, and Dynamics*, Vol. 33, No. 5, 2010, pp. 1429–1441. https://doi.org/10.2514/1.48475.

[16] Tordesillas, J., and How, J. P., "MADER: Trajectory Planner in Multiagent and Dynamic Environments," *IEEE Transactions on Robotics*, Vol. 38, No. 1, 2021, pp. 463–476. https://doi.org/10.1109/TRO.2021.3080235.

[17] Wu, S., Chen, G., Shi, M., and Alonso-Mora, J., "Decentralized Multi-Agent Trajectory Planning in Dynamic Environments with Spatiotemporal Occupancy Grid Maps," *arXiv preprint arXiv:2404.15602*, 2024.

[18] Vinod, A. P., Safaoui, S., Summers, T. H., Yoshikawa, N., and Di Cairano, S., "Decentralized, Safe, Multiagent Motion Planning for Drones Under Uncertainty via Filtered Reinforcement Learning," *IEEE Transactions on Control Systems Technology*, 2024. https://doi.org/10.1109/TCST.2024.3123456.

[19] Lindqvist, B., Sopasakis, P., and Nikolakopoulos, G., "A scalable distributed collision avoidance scheme for multi-agent UAV systems," *2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, IEEE, 2021, pp. 9212–9218. https://doi.org/10.1109/IROS51168.2021.9636293.

[20] Taye, A. G., Valenti, R., Rajhans, A., Mavrommati, A., Mosterman, P. J., and Wei, P., "Safe and Scalable Real-Time Trajectory Planning Framework for Urban Air Mobility," *Journal of Aerospace Information Systems*, 2024, pp. 1–10. https://doi.org/10.2514/1.I011381.

[21] Marshall, J. A., Sun, W., and L'Afflitto, A., "A survey of guidance, navigation, and control systems for autonomous multi-rotor small unmanned aerial systems," *Annual Reviews in control*, Vol. 52, 2021, pp. 390–427. https://doi.org/10.1016/j.arcontrol.2021.10.013.

[22] Rienecker, H., Hildebrand, V., and Pfifer, H., "Energy optimal 3D flight path planning for unmanned aerial vehicle in urban environments," *CEAS Aeronautical Journal*, Vol. 14, No. 3, 2023, pp. 621–636. https://doi.org/10.1007/s13272-023-00666-x.

[23] Ebert, C., Weiss, J., Uijt de Haag, M., Ruwisch, C., and Silvestre, F. J., "Trajectory Planning in Windy Urban Environment Using Gappy Proper Orthogonal Decomposition for Wind Estimates," *AIAA Journal*, Vol. 61, No. 6, 2023, pp. 2640–2651. https://doi.org/10.2514/1.J062049.

[24] Baskar, D., and Gorodetsky, A., "A Simulated Wind-field Dataset for Testing Energy Efficient Path-Planning Algorithms for UAVs in Urban Environment," *AIAA AVIATION 2020 FORUM*, 2020, p. 2920. https://doi.org/10.2514/6.2020-2920.

[25] Al-Sabban, W. H., Gonzalez, L. F., and Smith, R. N., "Wind-energy based path planning for unmanned aerial vehicles using Markov decision processes," *2013 IEEE International Conference on Robotics and Automation*, IEEE, 2013, pp. 784–789. https://doi.org/10.1109/ICRA.2013.6630662.

[26] Banerjee, P., and Bradner, K., "Energy-Optimized Path Planning for UAS in Varying Winds Via Reinforcement Learning," *AIAA AVIATION Forum and ASCEND 2024*, 2024, p. 4545. https://doi.org/10.2514/6.2024-4545.

[27] Souto, A., Alfaia, R., Cardoso, E., Araújo, J., and Francês, C., "UAV path planning optimization strategy: Considerations of urban morphology, microclimate, and energy efficiency using Q-learning algorithm," *Drones*, Vol. 7, No. 2, 2023, p. 123. https://doi.org/10.3390/drones7020123.

[28] Corbetta, M., and Kulkarni, C. S., "An approach for uncertainty quantification and management of unmanned aerial vehicle health," *Annual PHM Society Conference*, 2019. https://doi.org/10.36001/phmconf.2019.v11i1.847.

[29] Quinones-Grueiro, M., Biswas, G., Ahmed, I., Darrah, T., and Kulkarni, C., "Online decision making and path planning framework for safe operation of unmanned aerial vehicles in urban scenarios," *International Journal of Prognostics and Health Management*, Vol. 12, No. 3, 2021. https://doi.org/10.36001/ijphm.2021.v12i3.2953.

[30] Schumann, J., Kulkarni, C., Lowry, M., Bajwa, A., Teubert, C., and Watkins, J., "Prognostics for Autonomous Electric-Propulsion Aircraft," *International Journal of Prognostics And Health Management*, Vol. 12, No. 3, 2021. https://doi.org/10.36001/ijphm.2021.v12i3.2940.

[31] Taye, A., Thompson, E. L., Wei, P., Bonin, T., and Jones, J. C., "Probabilistic Evaluation for Flight Mission Feasibility of a Small Octocopter in the Presence of Wind," *AIAA AVIATION 2023 Forum*, 2023, p. 3964. https://doi.org/10.2514/6.2023-3964.

[32] Taye, A. G., and Wei, P., "Flight Mission Feasibility Assessment of Urban Air Mobility Operations under Battery Energy Constraint," *AIAA SCITECH 2024 Forum*, 2024, p. 0532. https://doi.org/10.2514/6.2024-0532.

[33] Daigle, M., and Kulkarni, C. S., "Electrochemistry-based battery modeling for prognostics," *Annual Conference of the PHM Society*, Vol. 5, 2013. https://doi.org/10.36001/phmconf.2013.v5i1.2252.

[34] Sudbury, A. W., and Hutchinson, E. B., "A cost analysis of amazon prime air (drone delivery)," *Journal for Economic Educators*, Vol. 16, No. 1, 2016, pp. 1–12.

[35] Taye, A. G., Bertram, J., Fan, C., and Wei, P., "Reachability based Online Safety Verification for High-Density Urban Air Mobility Trajectory Planning," *AIAA AVIATION 2022 Forum*, 2022, p. 3542. https://doi.org/10.2514/6.2022-3542.

[36] Corbetta, M., Banerjee, P., Okolo, W., Gorospe, G., and Luchinsky, D. G., "Real-time uav trajectory prediction for safety monitoring in low-altitude airspace," *AIAA AVIATION 2019 forum*, 2019, p. 3514. https://doi.org/10.2514/6.2019-3514.

[37] Teubert, C., Jarvis, K., Corbetta, M., Kulkarni, C., and Daigle, M., "ProgPy: Python Packages for Prognostics and Health Management of Engineering Systems," *Journal of Open Source Software*, Vol. 8, No. 87, 2023, p. 5099. https://doi.org/10.21105/joss.05099.

[38] Wan, E. A., and Van Der Merwe, R., "The unscented Kalman filter," *Kalman filtering and neural networks*, 2001, pp. 221–280. https://doi.org/10.1002/0471221546.ch7.

[39] Ekwaro-Osire, S., Gonçalves, A. C., and Alemayehu, F. M., *Probabilistic Prognostics and Health Management: A Brief Summary*, 1st ed., Springer, Cham, 2017, pp. 3–7. https://doi.org/10.1007/978-3-319-55852-3_1.

[40] Ahmed, I., Quinones-Grueiro, M., and Biswas, G., "A high-fidelity simulation test-bed for fault-tolerant octo-rotor control using reinforcement learning," *2022 IEEE/AIAA 41st Digital Avionics Systems Conference (DASC)*, IEEE, 2022, pp. 1–10. https://doi.org/10.1109/DASC55683.2022.9925862.

[41] Baskar, D., and Gorodetsky, A., "A Simulated Wind-field Dataset for Testing Energy Efficient Path-Planning Algorithms for UAVs in Urban Environment," , 2019. https://doi.org/10.7302/pdcv-0x63.