

A Dual-Approach Framework for eVTOL Climb Noise Mitigation

** All authors contributed equally to this research paper. Order listed alphabetically by surname.*

Matthias Casanova¹, Ethan Chien², Vismay Prasad³, Tuhina Samaddar⁴, Anya Talwar⁵, Landau Tzou⁶

¹Saratoga High School

²Greensboro Day School

³Saint John's High School

⁴Adrian Wilcox High School

⁵St. Francis High School

⁶Saint Francis High School

NASA STI Program . . . in Profile

Since its founding, NASA has been dedicated to the advancement of aeronautics and space science. The NASA scientific and technical information (STI) program plays a key part in helping NASA maintain this important role.

The NASA STI Program operates under the auspices of the Agency Chief Information Officer. It collects, organizes, provides for archiving, and disseminates NASA's STI. The NASA STI Program provides access to the NASA Aeronautics and Space Database and its public interface, the NASA Technical Report Server, thus providing one of the largest collections of aeronautical and space science STI in the world. Results are published in both non-NASA channels and by NASA in the NASA STI Report Series, which includes the following report types:

- **TECHNICAL PUBLICATION.** Reports of completed research or a major significant phase of research that present the results of NASA programs and include extensive data or theoretical analysis. Includes compilations of significant scientific and technical data and information deemed to be of continuing reference value. NASA counterpart of peer-reviewed formal professional papers, but having less stringent limitations on manuscript length and extent of graphic presentations.
- **TECHNICAL MEMORANDUM.** Scientific and technical findings that are preliminary or of specialized interest, e.g., quick release reports, working papers, and bibliographies that contain minimal annotation. Does not contain extensive analysis.
- **CONTRACTOR REPORT.** Scientific and technical findings by NASA-sponsored contractors and grantees.


- **CONFERENCE PUBLICATION.** Collected papers from scientific and technical conferences, symposia, seminars, or other meetings sponsored or co-sponsored by NASA.
- **SPECIAL PUBLICATION.** Scientific, technical, or historical information from NASA programs, projects, and missions, often concerned with subjects having substantial public interest.
- **TECHNICAL TRANSLATION.** English-language translations of foreign scientific and technical material pertinent to NASA's mission.

Specialized services also include creating custom thesauri, building customized databases, and organizing and publishing research results.

For more information about the NASA STI Program, see the following:

- Access the NASA STI program home page at <http://www.sti.nasa.gov>
- E-mail your question to help@sti.nasa.gov
- Fax your question to the NASA STI Information Desk at 443-757-5803
- Phone the NASA STI Information Desk at 443-757-5802
- Write to:
STI Information Desk
NASA Center for Aerospace Information
7115 Standard Drive
Hanover, MD 21076-1320

NASA/TM-2025AAABBBBCCC



A Dual-Approach Framework for eVTOL Climb Noise Mitigation

** All authors contributed equally to this research paper. Order listed alphabetically by surname.*

Matthias Casanova¹, Ethan Chien², Vismay Prasad³, Tuhina Samaddar⁴, Anya Talwar⁵, Landau Tzou⁶

¹Saratoga High School

²Greensboro Day School

³Saint John's High School

⁴Adrian Wilcox High School

⁵St. Francis High School

⁶Saint Francis High School

National Aeronautics and
Space Administration

Ames Research Center
Moffett Field, California 94035

August 2025

Acknowledgments

The successful realization of this research undertaking was made possible by the invaluable guidance and expertise of our mentors, Priyank Pradeep and Lindsay Stevens. We also extend our sincere gratitude to the NASA Ames Aviation Systems Division for granting us the opportunity for this project. In addition, we would also like to thank Jinhua Li and Hokkwan Ng for their work on AirNoiseUAM, which supported our noise modeling efforts, and Joseph Silva for developing and maintaining the Fe³ simulation framework.

The use of trademarks or names of manufacturers in this report is for accurate reporting and does not constitute an official endorsement, either expressed or implied, of such products or manufacturers by the National Aeronautics and Space Administration.

Available from:

NASA Center for AeroSpace Information
7115 Standard Drive
Hanover, MD 21076-1320

National Technical Information Service
5301 Shawnee Road
Alexandria, VA 22312

Available electronically at <http://www.sti.nasa.gov>

A Dual-Approach Framework for eVTOL Climb Noise Mitigation

Summary

Technological advances in Urban Air Mobility (UAM) will bring new aircraft to the skies above metropolitan regions. With noise pollution emerging as a major barrier to public acceptance, mitigating the sound produced by electric vertical takeoff and landing (eVTOL) vehicles is essential to the future viability of UAM. This paper focuses on the climb phase of eVTOL operations, one of the loudest and most prolonged noise-generating segments of flight. We present two viable approaches for generating and evaluating climb trajectories with respect to both acoustic impact and energy use. The first uses direct collocation via the PSOPT (Problem Solving for Optimal Control) framework, an open-source software package for optimal control problems, to generate climb trajectories of a commercial quadrotor at varying gradients. These trajectories are then evaluated for Sound Exposure Level (SEL) and energy consumption using a simplified model in lieu of full-scale simulation tools; we report separate minima for noise exposure and energy consumption rather than a weighted multi-objective optimum. The second approach, developed in parallel by our team, outlines a deep reinforcement learning (DRL) framework to explore climb planning from a data-driven perspective. In this setup, a DRL agent, using Q-learning algorithms, is designed to adjust climb profiles based on feedback from SEL and energy metrics, offering a data-driven complement to the model-based approach. Together, these methods lay the groundwork for a modular, scalable framework to support noise- and energy-aware trajectory design in future eVTOL operations. While our paper outlines both methodologies, only preliminary results are presented for the optimal control approach, with DRL training and multi-objective weighting left for future work.

Keywords: Electric Vertical Takeoff and Landing (eVTOL), Climb Trajectory Optimization, Noise Mitigation, Sound Exposure Level (SEL), Energy Efficiency, Direct Collocation, Double Deep Q-Network (DDQN), PSOPT

Contents

1	Nomenclature	3
2	Introduction	4
3	Background and Motivation	5
3.1	Background	5
3.2	Motivation	6
4	Problem Formulation	6
5	Assumptions	7
6	Optimal Control Framework	8
6.1	Optimal Control Overview	8
6.1.1	Direct Collocation Method	8
6.2	Vehicle Dynamics Model	9
6.2.1	Simplified Dynamic Formulation	9
6.2.2	Power Consumption Modeling	11
6.3	Acoustic Emission Modeling	12
6.3.1	SEL Calculation	12
6.3.2	Acoustic Model Parameters	13
6.3.3	Modeling Scope and Limitations	13
6.4	Trajectory Performance Evaluation	14
7	Deep Reinforcement Learning Framework	14
7.1	DRL Methodology	14
7.1.1	DRL Approach Overview	14
7.1.2	Simulation Environment	15
7.1.3	Reinforcement Learning	17
7.1.4	Double Deep Q-Learning Network (DDQN)	17
7.2	Reinforcement Learning Framework	18
7.2.1	Environment	18
7.2.2	State Space	18
7.2.3	Action Space	19
7.2.4	Reward Function	19
8	Results	21
9	Conclusions	22
10	Future Work	23

1 Nomenclature

Table 1.—Description of Variables Used in Vehicle Dynamics and Reinforcement Learning Framework. SI units are used throughout; acoustic levels are reported in dB re 20 μ Pa, unless noted.

Symbol	Description
λ	Latitude of the aircraft
τ	Longitude of the aircraft
h	Altitude above mean sea level of the aircraft
m	Mass of the aircraft
V	True airspeed of the aircraft
V_l	Lateral component of true airspeed
V_v	Vertical component of true airspeed
V_g	Groundspeed of the aircraft
γ	Aerodynamic flight path angle of the aircraft
ψ	Heading angle of the aircraft
χ	Course angle of the aircraft
D	Parasite drag on the aircraft
T	Net thrust
θ	Rotor tip-path-plane pitch angle
ϕ	Rotor tip-path-plane roll angle
κ	Induced power factor
ω	Rotational speed of the rotor blades
σ	Thrust-weighted solidity ratio
$C_{d,\text{mean}}$	Mean blade drag coefficient
A_{rotor}	Rotor disk area
R	Radius of the rotor
v_i	Rotor induced velocity during forward flight
T_{rotor}	Thrust produced by an isolated rotor
C_T	Rotor thrust coefficient
P_{required}	Power required for flight
ρ	Air density
R_{Earth}	Radius of the Earth
w_e	East component of wind velocity
w_n	North component of wind velocity
w_v	Vertical component of wind velocity
T_0	Reference time, 1 second
$p(t)$	Instantaneous sound pressure
p_{ref}	Reference sound pressure
t_1, t_2	Start and end times of the acoustic event
N	Total number of discrete trajectory nodes
p_i^2	Instantaneous squared sound pressure at node i

Continued on next page

Table 1 – continued from previous page

Symbol	Description
Δt_i	Time interval associated with node i
d	Three-dimensional Euclidean distance between the eVTOL and listener
h_{eVTOL}	Instantaneous eVTOL altitude
h_{listener}	Listener’s altitude
$L_{s,1m}$	Source sound pressure level at 1 meter
T_{ac}	Average total thrust of the eVTOL
$L_{\text{ref},1m}$	Reference sound pressure level at 1 meter
n_T	Constant factor
T_{ref}	Reference thrust
P_{avg}	Average power
V_{airspeed}	Average airspeed of the eVTOL
L_p	Received instantaneous sound pressure level at the listener’s location
d_{ref}	Reference distance, 1 meter
ΔL_{GR}	Ground reflection factor
$L_{p,i}$	Instantaneous sound pressure level at node i
a_t	Action vector at time t
s_t	State vector at time t
r_t	Intermediate reward at time t
R_T	Terminal reward
$Q(s, a)$	Q-value for state–action pair
$V(s)$	Value of being in state s
π	Policy function
Δd_{OD}	Distance-to-goal from origin to destination per timestep
ΔE_t	Energy used per timestep
ΔE_{total}	Total energy used
SEL_t	Sound Exposure Level at time t
w_1, w_2, \dots, w_5	Reward weighting coefficients
ϵ -greedy	Exploration–exploitation decay mechanism

2 Introduction

Urban Air Mobility (UAM) and Advanced Air Mobility (AAM) are emerging transportation systems designed to alleviate surface congestion by enabling short-range, low-altitude transportation using electric vertical takeoff and landing (eVTOL) aircraft. These vehicles are intended for applications such as air taxis, emergency response, and intermodal connectivity in dense urban environments (Ref. 1, 2). They rely on compact platforms like quadrotors capable of vertical takeoff and landing in constrained settings.

Despite their promise, UAM systems face several operational challenges. A primary concern is community noise, which poses a barrier to public acceptance and regulatory approval (Ref. 3). Unlike high-altitude commercial flights, eVTOL aircraft operate near the ground and frequently over populated areas, increasing acoustic exposure. Additionally, battery limitations constrain

range and flight profiles (Ref. 4).

This paper focuses on investigating strategies for mitigating noise during the climb phase, where eVTOL aircraft operate at high thrust and low altitude, producing significant acoustic output. Sound Exposure Level (SEL), a standard metric that integrates noise intensity and duration, is used throughout this study to quantify acoustic impact (Ref. 5–7). While noise along fixed trajectories has been studied, few works have explored how trajectory-level adjustments, such as climb angle, might reduce noise.

To address this, we evaluate two complementary approaches. The first employs direct collocation via the PSOPT optimal control framework to generate climb trajectories for a reference quadrotor, assessed with a simplified noise model. The second leverages a Double Deep Q-Network (DDQN) reinforcement learning agent to learn noise-aware strategies through feedback from noise and energy metrics.

The remainder of this paper is structured as follows. Section 3 reviews background and motivation. Section 4 states the problem formulation. Section 5 summarizes modeling assumptions. Section 6 develops the optimal control framework, including vehicle dynamics, power modeling, acoustic emission modeling, and the trajectory evaluation workflow. Section 7 presents the deep reinforcement learning framework, covering the simulation environment and agent design. Section 8 reports results. Section 9 provides conclusions. Section 10 outlines future work.

3 Background and Motivation

3.1 Background

UAM systems aim to offer efficient, low-emission aerial transportation in metropolitan regions. By bypassing surface congestion, these systems could improve urban mobility and reduce travel times. At the core of UAM are eVTOL aircraft, designed for vertical operations in tight urban spaces. These vehicles are expected to operate at relatively low speeds and altitudes, often in corridors shared with other air traffic, infrastructure, and environmental constraints.

However, the deployment of eVTOLs introduces unique acoustic challenges. Unlike conventional aircraft that rapidly ascend to cruise altitudes, eVTOLs remain within urban airspace for a longer portion of their flights. As a result, noise emissions become more relevant to communities directly below flight paths. Rotorcraft, especially multirotor eVTOLs, produce nonstationary, broadband noise that varies with operating conditions. During low-altitude segments, such as takeoff, climb, and landing, this noise becomes particularly prominent (Ref. 8).

Technical factors such as rotor configuration, tip speed, blade loading, and control surface activity influence both the frequency content and magnitude of acoustic emissions (Ref. 9). In particular, blade-vortex interactions and rapid thrust changes can generate tonal components that are more perceptible and objectionable to human listeners (Ref. 10). The growing diversity in vehicle designs, ranging from coaxial rotors to tilt-wing architectures, further complicates noise prediction and certification efforts.

In response, regulatory bodies are working to develop noise certification metrics that reflect the distinct operational patterns of UAM systems, with a focus on frequent, low-altitude flights in noise-sensitive areas (Ref. 11). However, most existing tools evaluate noise along static, predefined trajectories. There is limited integration between flight dynamics, vehicle control, and acoustic modeling in current planning frameworks. As a result, there is a growing interest in dynamic

trajectory design that can respond to noise constraints in real time.

Early studies such as Pradeep et al. (Ref. 12) explored the energy and time implications of different climb profiles for multirotor vehicles. These efforts provide a basis for examining how trajectory-level control could influence acoustic outcomes in urban airspace. Building on this direction, our work focuses on using physically grounded models to analyze the potential of climb trajectory design as a tool for noise mitigation.

3.2 Motivation

Community noise has emerged as one of the most persistent obstacles to the widespread adoption of UAM services (Ref. 13). Aircraft noise has been linked to a range of negative health outcomes, including elevated stress, cardiovascular issues, and impaired sleep. In children, chronic exposure is associated with reduced cognitive performance and behavioral disturbances (Ref. 14). These effects are amplified in dense urban settings, where repeated exposure from multiple flights can accumulate over time.

Unlike conventional airports, future vertiports may be embedded directly into neighborhoods, atop office buildings, hospitals, or transit hubs. This proximity increases the importance of understanding how different flight phases contribute to perceived noise. Among these phases, climb is particularly consequential. It involves high power output and relatively slow ascent, which extends the duration of noise exposure near the departure point.

Current trajectory optimization methods often emphasize objectives like energy efficiency or travel time. While these are critical to operational feasibility, they may overlook how small changes in climb angle or acceleration affect the spatial and temporal distribution of noise. By incorporating noise considerations into trajectory design, UAM operators could reduce community impact without significant performance tradeoffs.

In this study, we evaluate whether noise-aware climb profiles can be developed using either optimal control or reinforcement learning approaches. We simulate a commercially scaled quadrotor (Ref. 15) using simplified but consistent models of aerodynamics, power consumption, and noise propagation. While this paper emphasizes early results from the optimal control approach, the broader goal is to lay the foundation for adaptive trajectory planning methods that support both operational efficiency and public acceptability.

4 Problem Formulation

A fundamental challenge in UAM deployment is minimizing noise at the point of departure. Vertiports may be located atop hospitals, schools, or residential buildings, sites with low tolerance for acoustic disturbance. Reducing sound exposure near takeoff locations is therefore crucial for public trust and scalable operations.

This work focuses on the climb phase, where the vehicle transitions from liftoff to cruise. During this period, it operates at high power while in close proximity to ground observers. The trajectory geometry directly influences the intensity and duration of noise received. Figure 1 illustrates how different climb angles alter acoustic impact for a microphone located at the vertiport.




Figure 1.—Flight phase terminology for Urban Air Mobility (UAM) operations. The climb phase is characterized by flight path angle γ , with three example angles illustrated: steep ($\gamma = 60^\circ$), moderate ($\gamma = 45^\circ$), and shallow ($\gamma = 30^\circ$).

Noise exposure is nonlinear and sensitive to both trajectory and vehicle characteristics. Even small changes in flight path can significantly alter the acoustic footprint. This work supports the goals of NASA’s Advanced Air Mobility (AAM) Project by investigating trajectory-based noise mitigation strategies that promote community integration and safe, acceptable urban flight (Ref. 16).

5 Assumptions

To enable simpler modeling and reduce computational complexity, we introduce several assumptions regarding our system. These assumptions help guide the formulation of the simplified dynamics model used in our trajectory generation.

From an environmental standpoint, we include a constant wind component during the climb phase, ignoring wind variability with altitude due to minimal altitude discrepancies during the climb. Specifically, we assume a steady 5 m/s eastward wind based on average wind patterns observed in San Francisco (Ref. 17, 18).

Operationally, the climb begins at an altitude of 15 m above ground level (AGL), to account for the vehicle’s transition phase, and ends at 488 m, a typical cruise altitude for urban air taxis. The climb is modeled as a straight ascent with a constant heading of 180° , assuming no horizontal maneuvering during the climb. We also assume a constant lateral airspeed of 30 m/s, as established in prior UAM literature (Ref. 19). This ensures that as the vehicle increases its climb angle (γ), the power consumption adjusts accordingly. Additionally, for our deep reinforcement learning setup, we adopt minimum energy consumption (Δe) as the baseline criterion, enabling a direct comparison

between energy-optimal climb trajectories and noise-minimizing climb profiles.

For simulation setup, the trajectory origin is fixed above Golden Gate Park in San Francisco, chosen to provide a consistent geographic reference, while also simplifying altitude measurements, since MSL (Mean Sea Level) and AGL measurements are effectively equal at Golden Gate Park due to its proximity to sea level. Lastly, we model the vehicle as the quadrotor seen in Figure 2, a well-characterized eVTOL platform with detailed performance specifications, seen in Table 2, publicly available in the UAM research community (Ref. 15).




Figure 2.—Quadrotor eVTOL (Ref. 15)

Table 2.—Quadrotor Parameters Used in Dynamics and Power Models (Ref. 12)

Parameter	Symbol	Value
Rotor radius	R	4.0 m
Rotor disk area	A_{rotor}	50.26 m ²
Mass	m	2940 kg
Fuselage drag correction factor	F_p	0.97
Thrust coefficient	C_T	0.0055
Maximum velocity	V_{max}	56 m/s

6 Optimal Control Framework

6.1 Optimal Control Overview

6.1.1 Direct Collocation Method

To generate climb trajectories that are both feasible and energy-efficient, we frame the task at hand as an optimal control problem. This allows us to find control inputs, such as thrust and blade pitch, that guide the aircraft from its starting point to a target altitude while minimizing energy use and staying within physical boundaries. We use direct collocation, a numerical method that is well-suited for problems involving nonlinear dynamics and multiple constraints, such as those

present in our system.

In direct collocation, the continuous flight path is broken up into a series of time steps, called collocation points. At each point, the aircraft’s motion is approximated using polynomial curves, and the laws of motion, expressed as differential equations, are applied as algebraic constraints. This process converts the continuous-time problem into a nonlinear programming (NLP) problem. A detailed explanation of direct collocation can be found in Kelly’s article (Ref. 20).

We implement this approach using PSOPT (Problem Solving for Optimal Control), an open-source software that handles the transcription of system dynamics and constraints. The resulting NLP is solved using IPOPT (Interior Point OPTimizer), a numerical solver for large-scale constrained optimization problems.

6.2 Vehicle Dynamics Model

6.2.1 Simplified Dynamic Formulation

To aid in understanding the variables and reference frames used in our formulation, Figures 3 and 4 illustrate the key geometry and pitch-plane variables applied in the vehicle dynamics model. Figure 3 shows the navigation frame geometry, including wind, airspeed, and ground speed vectors, along with heading and track angles. Figure 4 depicts the pitch-plane variables, highlighting the aircraft reference point and climb-related angles. These definitions establish the notation applied in Equations 1–14.

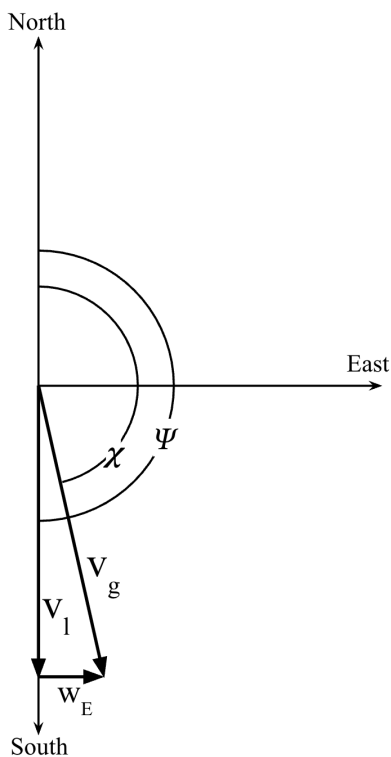


Figure 3.—Navigation frame geometry showing wind vector W_E , airspeed vector V_I , ground speed vector V_g , and angles ψ (heading) and χ (track).




Figure 4.—Variable definitions in the pitch plane. The purple dot denotes the aircraft reference point. A zero bank angle ($\phi = 0$) is assumed, meaning the roll axis (not shown) is perpendicular to the pitch plane.

To implement our trajectory optimization in PSOPT, we require a set of nonlinear dynamic equations that describe the motion of the quadrotor. These equations must define how the system’s state variables evolve over time in response to specific control inputs. Based on the assumptions outlined in the previous section, we begin with the full 6-axis dynamic model of the commercial quadrotor platform as described in Pradeep et al.’s study and seen in equations 1-7 (Ref. 12). We then simplify this model to suit our specific use case, vertical climb under fixed heading, eastward wind, and constant lateral velocity.

$$\frac{dV_l}{dt} = \frac{T \cos \phi \sin \theta - D \cos \gamma}{m} \quad (1)$$

$$\frac{dV_v}{dt} = \frac{T \cos \phi \cos \theta - D \sin \gamma - mg}{m} \quad (2)$$

$$\frac{d\psi}{dt} = \frac{T \sin \phi}{mV_l} \quad (3)$$

$$\frac{d\lambda}{dt} = \frac{V_l \cos \psi + w_n}{R_{\text{Earth}} + h} = \frac{V_g \cos \chi}{R_{\text{Earth}} + h} \quad (4)$$

$$\frac{d\tau}{dt} = \frac{V_l \sin \psi + w_e}{(R_{\text{Earth}} + h) \cos \lambda} = \frac{V_g \sin \chi}{(R_{\text{Earth}} + h) \cos \lambda} \quad (5)$$

$$\frac{dh}{dt} = V_v + w_v \quad (6)$$

$$\tan \gamma = \frac{V_v}{V_l} \quad (7)$$

The resulting differential equations 8-14, shown below, represent the vehicle’s longitudinal and

vertical motion and serve as the system dynamics passed to the direct collocation solver.

$$\frac{dV_l}{dt} = 0 \quad (8)$$

$$\frac{dV_v}{dt} = \frac{T \cos \theta - D \sin \gamma - mg}{m} \quad (9)$$

$$\frac{d\psi}{dt} = 0 \quad (10)$$

$$\frac{d\lambda}{dt} = \frac{30 \cdot \cos(180^\circ) + 0}{R_{\text{Earth}} + h} = \frac{-30}{R_{\text{Earth}} + h} \quad (11)$$

$$\frac{d\tau}{dt} = \frac{30 \cdot \sin(180^\circ) + 5}{(R_{\text{Earth}} + h) \cos(\lambda)} = \frac{5}{(R_{\text{Earth}} + h) \cos(\lambda)} \quad (12)$$

$$\frac{dh}{dt} = V_v + 0 = V_v \quad (13)$$

$$\tan \gamma = \frac{V_v}{30} \quad (14)$$

To simplify the aircraft dynamics for trajectory solving in PSOPT, we apply assumptions aligned with the climb phase of the quadrotor. The heading angle is fixed at $\psi = 180^\circ$, corresponding to due south in aviation coordinates, which yields $\cos \psi = -1$ and $\sin \psi = 0$. These values simplify the geodetic equations for latitude and longitude. The lateral (horizontal) velocity is held constant at $V_l = 30$ m/s, which implies $\frac{dV_l}{dt} = 0$. We assume zero roll, $\phi = 0$, and no yawing motion, so $\frac{d\psi}{dt} = 0$.

The wind field is assumed to have only an eastward component of 5 m/s, i.e., $(w_n, w_e, w_v) = (0, 5, 0)$. This eliminates the vertical and northward wind terms. As a result, the latitude rate reduces to a constant, $\frac{d\lambda}{dt} = \frac{-30}{(R_{\text{Earth}} + h)}$, and the longitude rate simplifies to $\frac{d\tau}{dt} = \frac{5}{(R_{\text{Earth}} + h) \cos \lambda}$ due to the combined effects of southward motion and eastward wind. The altitude rate is given by $\frac{dh}{dt} = V_v$, and the climb angle is expressed as $\tan \gamma = \frac{V_v}{30}$. These simplifications preserve the essential dynamics for evaluating climb energy and noise while allowing for ease in trajectory solving.

6.2.2 Power Consumption Modeling

We adopt the total power model for rotorcraft flight used in (Ref. 12), which is derived from classical formulations of helicopter aerodynamics (Ref. 15, 21). This model is well-suited for integration into our trajectory optimization framework, which evaluates state-dependent power at each collocation node. The total required power includes induced, parasitic, climb, and profile components:

$$P_{\text{required}} = P_{\text{induced}} + P_{\text{parasite}} + P_{\text{climb}} + P_{\text{profile}} \quad (15)$$

Following (Ref. 12), we compute the total power as:

$$P_{\text{required}} = \kappa \sum_{n=1}^4 (T_{\text{rotor}} v_i)_n + TV \sin \alpha + \frac{\rho A_{\text{rotor}} (\omega R)^3 \sigma C_{d,\text{mean}} F_P}{8} \quad (16)$$

Here, v_i is the induced velocity of each rotor (solved numerically), α is the angle of attack, and σ , $C_{d,\text{mean}}$, and F_P represent blade-specific aerodynamic parameters. κ accounts for empirical

corrections to the induced power model. While this equation is taken from prior literature, we uniquely apply it within a PSOPT-based optimal control framework and focus specifically on climb profiles under fixed heading, roll, and lateral speed.

To support these calculations, we extract relevant physical characteristics of the quadrotor (e.g., rotor radius, mass) from the vehicle specification table (Table 2).

6.3 Acoustic Emission Modeling

To evaluate the noise generated by an eVTOL climb, a simplified noise model was developed as a separate Python script that uses the trajectory summary data produced by PSOPT. We used this due to limited time and resources, however, if the most accurate output were required, we would recommend using NASA’s AirNoiseUAM tool. The noise model estimates the SEL at a fixed listener location directly beneath the takeoff point, providing a preliminary approximation suitable for trend analysis and visualization.

6.3.1 SEL Calculation

The SEL is a measure of the total sound energy of a single acoustic event, normalized to a 1 s duration. The continuous formula for SEL is given by (Ref. 22):

$$SEL = 10 \log_{10} \left(\frac{1}{T_0} \int_{t_1}^{t_2} \frac{p^2(t)}{p_{\text{ref}}^2} dt \right) \quad (17)$$

For discrete implementation within the computational framework, this integral is approximated by a summation over trajectory nodes:

$$SEL = 10 \log_{10} \left(\sum_{i=1}^N \frac{p_i^2}{p_{\text{ref}}^2} \Delta t_i \right) \quad (18)$$

Here, p_i^2 is the instantaneous squared sound pressure at node i , $p_{\text{ref}}^2 = (20 \times 10^{-6} \text{ Pa})^2$ is the reference squared pressure, and Δt_i is the time interval associated with node i .

The instantaneous squared pressure, p_i^2 , for each node is determined through a sequence of calculations.

First, the three-dimensional distance (d) between the eVTOL and the listener at any given moment is computed. This is a Euclidean distance incorporating both horizontal ($d_{\text{horizontal}}$) and vertical separation:

$$d = \sqrt{(h_{\text{eVTOL}} - h_{\text{listener}})^2 + d_{\text{horizontal}}^2} \quad (19)$$

Here, h_{eVTOL} is the instantaneous eVTOL altitude, h_{listener} is the listener’s altitude (assumed 15.2 m), and $d_{\text{horizontal}}$ is the horizontal distance derived from the differences in longitude and latitude using a flat-Earth approximation.

The source sound pressure level at 1 meter ($L_{s,1m}$) is then estimated based on the eVTOL’s average (T_{avg}):

$$L_{s,1m} = L_{\text{ref},1m} + 10 \cdot n_T \cdot \log_{10} \left(\frac{T_{\text{avg}}}{T_{\text{ref}}} \right) \quad (20)$$

In this simplified model, T_{avg} is calculated as the average power required (P_{avg}) divided by the average airspeed ($V_{airspeed}$):

$$T_{avg} = \frac{P_{avg}}{V_{airspeed}} \quad (21)$$

Where P_{avg} is derived from energy consumption over the mission. This T_{avg} represents an equivalent total system thrust, serving as a proxy for the overall propulsive effort given the eVTOL's average power required. It is important to note that this is a simplification, as it does not distinguish between power consumed for pure propulsion versus power consumed to overcome induced or parasitic drag at a specific instant.

$L_{ref,1m}$, n_T , and T_{ref} are model parameters. Subsequently, the received instantaneous SPL (L_p) at the listener's location is calculated by accounting for spherical spreading loss and a constant ground reflection factor:

$$L_p = L_{s,1m} - \left(20 \log_{10} \left(\frac{d}{d_{ref}} \right) \right) + \Delta L_{GR} \quad (22)$$

Here, d_{ref} is the reference distance (1 m), and ΔL_{GR} is the ground reflection factor.

Finally, the instantaneous SPL is converted to squared pressure (p_i^2) using the reference pressure, to facilitate energy summation for SEL:

$$p_i^2 = p_{ref}^2 \cdot 10^{\left(\frac{L_{p,i}}{10} \right)} \quad (23)$$

This p_i^2 value is then used in the discrete SEL summation.

6.3.2 Acoustic Model Parameters

The noise model relies on the following key parameters:

- $L_{ref,1m} = 170.0$ (dB): Reference SPL at 1 meter from the source.
- $T_{ref} = 28.8$ (kN): Reference thrust for noise scaling (e.g., hover thrust).
- $n_T = 1.5$: Exponent for the thrust-noise relationship.
- $\Delta L_{GR} = 3.0$ (dB): Additive factor for ground reflection.

6.3.3 Modeling Scope and Limitations

It is critical to acknowledge the significant simplifications inherent in this noise model. The noise model provides extremely simplified estimates to offer a general preview of the trend of SEL vs. energy consumed, and relies on numerous simplifying assumptions. Consequently, these results are not suitable for precise acoustic certification or detailed environmental impact assessments. Tools like NASA's AirNoiseUAM would be a much better alternative for more sophisticated and accurate noise modeling.

Key limitations of the current model include the following: The eVTOL is treated as a single, omnidirectional noise source, ignoring complex noise generation from multiple rotors and the airframe (Point Source Assumption). The flight profile is assumed to be a simplified steady-state climb with constant airspeed and a linear flight path. Furthermore, a basic empirical power law relates

total thrust to noise, which does not account for detailed aerodynamic noise mechanisms. Sound attenuation due to atmospheric conditions is not modeled. Noise radiation is assumed uniform in all directions, neglecting the directional nature of rotor noise. Finally, a fixed value is used for ground reflection, assuming a perfectly reflective, flat surface.

6.4 Trajectory Performance Evaluation

To evaluate the noise impact and energy tradeoffs of various climb profiles, we construct a multi-stage solution pipeline that combines trajectory generation, noise modeling, and data analysis. The complete process is illustrated in Figure 5.

We begin by specifying a set of fixed climb angles and inputting initial conditions and vehicle performance parameters. For each climb angle, we use the PSOPT framework to generate feasible trajectories using direct collocation. Although PSOPT is capable of solving optimal control problems, in this study it is used to produce valid 4-D trajectories under predefined control inputs. PSOPT also outputs time histories of relevant state and control variables, including energy consumption.

These trajectories are then passed to a noise estimation module, which produces SEL along the flight path. While NASA’s AirNoiseUAM tool is conceptually intended for this purpose, preliminary results are obtained using a simplified internal noise model due to time constraints.

Finally, both energy consumption and SEL are passed to a data analysis step. Here, we construct a tradeoff analysis between noise and energy to assess the cost of each climb angle and identify a set of favorable profiles for further study.

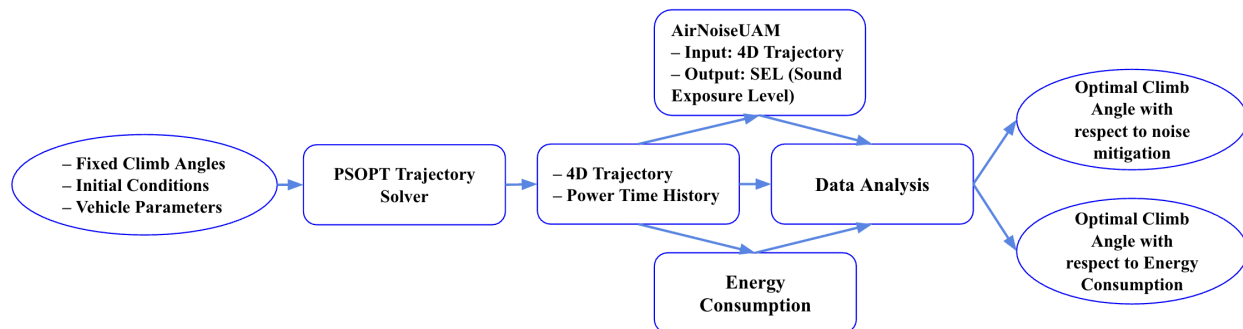


Figure 5.—Multi-stage evaluation of climb profiles using trajectory solver, noise modeling and data analysis

7 Deep Reinforcement Learning Framework

7.1 DRL Methodology

7.1.1 DRL Approach Overview

While optimal control generates noise-aware climb profiles via model-based transcriptions, Deep Reinforcement Learning (DRL) offers a data-driven alternative capable of discovering novel maneuver strategies without explicit linearization or collocation. In this work, we frame climb-trajectory optimization as a Markov decision process in which an agent selects discrete climb-angle adjustments to minimize a weighted sum of SEL and energy consumption.

7.1.2 Simulation Environment

Two simulation frameworks that utilize NASA’s Flexible Engine for fast-time evaluation of flight environments Fe^3 (Ref. 23) are used to model the climb phase of an electric Vertical takeoff and Landing (eVTOL) vehicle. Flight dynamics, flight kinematics, energy consumption, and motor speed have been incorporated into Fe^3 to simulate eVTOL behavior. The reinforcement learning (RL) agent learns a policy that optimizes for energy consumption in one simulation framework and learns a policy that minimizes noise in a second simulation framework.

This method uses the deep learning library PyTorch and the reinforcement learning framework TorchRL to train the agent to ultimately minimize noise exposure and energy consumption by adjusting the flight path angle (γ) during the initial climb phase. During simulation, Fe^3 generates and updates the simulation time and eVTOL state at a fixed time step of 0.5s. These states (eVTOL, simulation time, maneuver information) are stored in shared GPU memory accessible by both the Fe^3 and Python processes. Based on these observations, the Python process computes an adjustment to the angle of climb and writes the result to the shared GPU memory. Fe^3 then reads the new maneuver input, updates the vehicle state, and simulates the vehicle’s next state. Compute Unified Device Architecture (CUDA), NVIDIA’s parallel computing platform for GPU programming, is used for inter-process communication (IPC) alongside Unix-named first-in-first-out pipes (FIFO) to share the state data between the Fe^3 and Python process (Ref. 24). This communication loop continues until the agent reaches its target altitude of 488 m.

Named pipes are used to ensure proper data flow and avoid race conditions between the Fe^3 and Python processes. The Python process opens two pipes: one that it can only read from and one that it can only write to. The Fe^3 process opens the same pipes in opposite modes: one in read-only mode and one in write-only mode. During the main simulation loop, these pipes are used to relay status messages when a process is done reading or writing simulation data.

Simulation data is stored on a shared GPU memory, to which both the Fe^3 and the Python processes have access. The shared GPU allows both processes to read and write to the same block with minimal overhead. Additionally, the Python process stores other tensors on the GPU, including the weights for the neural network policy. Fe^3 stores memory on the GPU necessary for simulation, such as way-point data and wind data.

As shown in Figure 6, the energy-optimization configuration is solely focused on minimizing energy consumption throughout the climb. The Python agent accesses both previous and current simulation time and the previous and current eVTOL states to calculate the energy expended (ΔE_s) over that interval and determine the effectiveness of the last maneuver. Using this pair of points, the agent computes a new angle-of-climb (γ) maneuver that aims to minimize energy used in the next pair of points. The latest action is written back to GPU memory, then read by Fe^3 for the next step in the simulation.

The noise-optimization configuration, shown in Figure 7, extends upon the energy-optimization framework by integrating the AirNoiseUAM module (Ref. 7), a software tool that models the noise exposure of Urban Air Mobility (UAM) operations. At each simulation step Fe^3 and the Python process record both the previous and current eVTOL states. The AirNoiseUAM module accesses this segment information, as well as preprocessed Noise-Power Distance (NPD) data, to compute the SEL (Ref. 22) at the vertiport 1. The RL (Reinforcement Learning) agent receives both the current and prior eVTOL states and evaluates the noise impact resulting from the previous action using the SEL value at the vertiport. The Python process then calculates a new climb maneuver

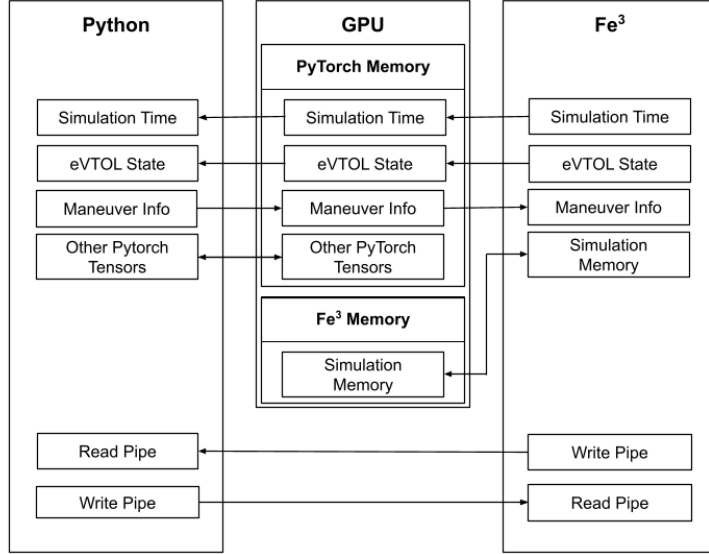


Figure 6.—Shared memory and inter-process communication between Python and Fe³ processes for energy-optimization configuration (Ref. 25).

that seeks to minimize the acoustic footprint while progressing towards the target altitude.

When used in tandem, the energy and noise optimization simulation frameworks enable a comprehensive exploration of tradeoffs between energy usage and community impact during the eVTOL climb phase. By independently training reinforcement learning under each objective (minimizing energy consumption and acoustic exposure-performance tradeoffs), areas of overlap or conflict between objectives can be quantified.

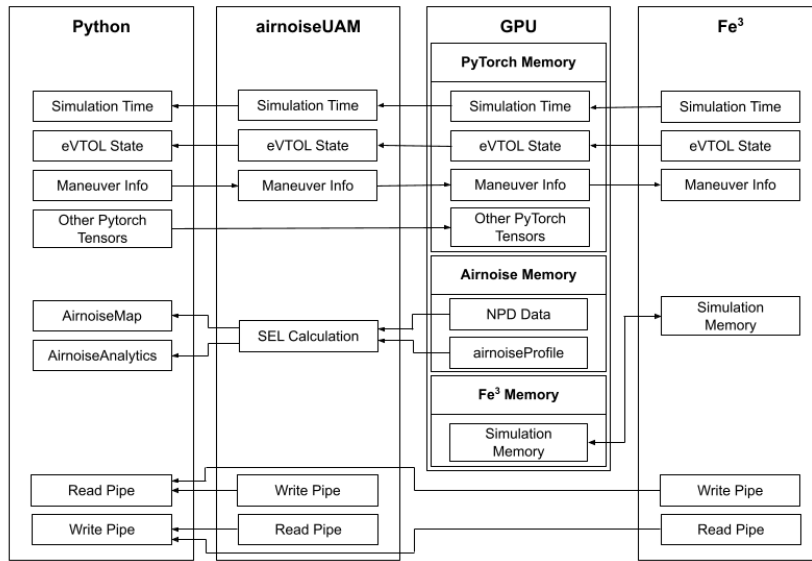


Figure 7.—Shared memory and inter-process communication between Python and Fe³ processes for noise-minimizing configuration (Ref. 25).

7.1.3 Reinforcement Learning

Reinforcement learning (RL) is a branch of machine learning that involves an agent taking actions in an environment to maximize a cumulative reward. An RL problem is defined by the components of the tuple (S, A, P, R, γ) , where γ is now defined as the discount factor. At each time step, t , the agent observes the current state ($S_t \in S$) and selects an action ($A_t \in A$) according to policy (π), which is a function mapping states to a probability distribution over actions. S and A represent all possible state and action spaces within the environment. Based on the state-action pair (S, A) , the state is then updated to S_{t+1} and the agent receives a reward (R_t). The environment's dynamics, defined by the transition probability $P(S_{t+1}|S_t, A_t)$, determine how the state progresses from S_t to S_{t+1} given action A_t . After the agent reaches state S_{t+1} , the reward function $R(R_{t+1}|S_t, A_t, S_{t+1})$ assigns the reward to the agent.

The agent's objective is to take actions according to the optimal policy (π^*) that maximizes the cumulative reward over time. This optimal policy is expressed as follows:

$$\pi^*(t) = \arg \max_{\pi} \mathbb{E} \left[\sum_{i=0}^{\infty} \gamma^i R_{t+i+1} \mid \pi \right] \quad (24)$$

A proper design of the reward function is critical, as it directs the agent's behavior by specifying desirable outcomes without telling the agent how to achieve the desired outcome (Ref. 26). The discount factor $\gamma \in (0, 1)$ determines the value of future rewards; when $\gamma \rightarrow 0$, the agent prioritizes immediate rewards and when $\gamma \rightarrow 1$, the agent emphasizes future rewards. Ultimately, reinforcement learning yields a policy that maximizes the long-term reward (Ref. 26).

There are two approaches used to train an agent to discover its optimal policy: value-based and policy-based methods. Value-based methods teach the agent to learn which state is more valuable through learning a value function that maps the expected value of being at that state. The state value function ($V_{\pi}(S_t)$), is the first representation of the value function, which expresses the expected return of acting according to the policy ($\mathbb{E}_{\pi} [R_{t+1} + \gamma R_{t+2} + \gamma^2 R_{t+3} + \dots \mid S_t]$) starting in state S_t . The action-value function or the second representation of the value function $Q_{\pi}(S_t, A_t)$ reflects the expected discounted reward when starting in state S_t , taking action A_t , and following the policy ($\mathbb{E}_{\pi} [R_{t+1} + \gamma R_{t+2} + \gamma^2 R_{t+3} + \dots \mid S_t]$). Contrastingly, policy-based methods ($\pi(S_t)=P(A_t|S_t)$) directly learn the optimal policy from the agent's interactions with the environment, without having to maintain a value function estimation. An optimal value function (V^* or Q^*) is reached when it meets the Bellman optimality equation (Ref. 27) in Equation 25.

$$\begin{aligned} V^*(S_t) &= \max_{A_t} \mathbb{E} [R_{t+1} + \gamma V^*(S_{t+1})] \\ Q^*(S_t, A_t) &= \mathbb{E} \left[R_{t+1} + \gamma \max_{A_{t+1}} Q^*(S_{t+1}, A_{t+1}) \right] \end{aligned} \quad (25)$$

7.1.4 Double Deep Q-Learning Network (DDQN)

A Double Deep Q-Network (DDQN) is an improved version of the traditional Deep Q-Network (DQN) used in Deep Reinforcement Learning (DRL). Like DQN, DDQN enables an agent to learn how to make optimal decisions in various situations, thereby maximizing rewards. The key improvement in DDQN is that it can accurately estimate the value of actions and reduce overestimation bias, which happens when the network consistently produces overly optimistic value estimates (Ref. 28).

DDQN is a temporal-difference (TD) learning method that updates value estimates based on the immediate reward and the estimated value of the next state, without waiting for an episode to complete. The main Q-learning update uses the “TD target” and measures the TD error to quantify the difference between the prediction and the actual outcome (Ref. 29).

DDQN handles this bias by using two neural networks instead of one: an online network and a target network. During training, the online network is used to take the best action for the next state, while the target network estimates the value of that action. By splitting the action selection and evaluation between two networks, DDQN gives a more stable and accurate value estimate and more robust policies. The target network is updated less frequently than the online network, helping to stabilize training.

To ensure the agent can explore various actions and avoid settling too early on a less effective choice, an epsilon-greedy policy is also used during training to balance exploration (trying new actions) and exploitation (choosing the best-known action). This approach allows the agent to select the highest-valued action a majority of the time, while occasionally choosing a random action with probability epsilon (ϵ). ϵ starts high to allow for exploration, and then is gradually reduced during training to allow the agent to exploit acquired knowledge (Ref. 29).

Additionally, DDQN uses experience replay, where as the agent interacts with the environment, the experiences (such as state, action, reward, next state) are stored in a tuple and sampled randomly for training the Q-network. This improves learning efficiency and mitigates forgetting of previous experiences.

7.2 Reinforcement Learning Framework

7.2.1 Environment

The simulation will create a relatively stable environment. The randomly generated situations will only adjust the climb angle(γ) at each 0.5s time step, maintaining constant heading. The climb angle will stay within a range of a minimum of 3°, and a maximum of 90°, ensuring the agent moves upwards and does not change lateral direction. The simulation will assume that there are no harsh weather conditions, with a consistent wind speed going eastward at 5 m/s derived from National Climatic Data Center observations, with altitude-dependent wind variations excluded from this preliminary approach (Ref. 30) and no flight path obstacles, allowing for simple optimal trajectories for noise and energy during the climb phase. The eVTOL will need to reach a cruise altitude of 488 m going at a constant lateral speed of 30 m/s and at that point the episode ends.

7.2.2 State Space

Clearly defining the state space (S_t) is crucial because it determines what information the agent uses to make decisions. A poorly designed state space can lead to ineffective learning or unsafe behavior (Ref. 29). The state space for our drone takeoff and landing simulation captures all the essential variables that describe the eVTOL’s situation at any given time step. This includes the climb angle (γ_t), as well as energy consumption(E_t), distance moved towards the destination (Δd_{OD}), and SEL value (SEL_t).

7.2.3 Action Space

The action space defines all possible moves or decisions an agent can take at any given state. To get the optimal action using the Q-function, the action space needs to be discretized such that the number of output neurons matches the number of actions (Ref. 29). The bounds for our action space are $\gamma_t \in [-30, 30]$ where γ_t is segmented at intervals of 1 degree. This means that at each timestep, the maximum by which γ can be adjusted is 30°

7.2.4 Reward Function

The reinforcement learning framework uses both intermediate and terminal rewards to get the eVTOL to have an energy-efficient and minimal sound level exposure value during its trajectory. The reward function is therefore split into two primary optimization objectives: energy efficiency and noise mitigation.

Intermediate rewards are measured and evaluated at every time step, t , to ensure that the agent can efficiently maintain its trajectory while prioritizing energy efficiency and noise mitigation. Terminal rewards are calculated at the end of an episode to allow the RL agent to reflect on the success of their actions in terms of the total energy used and the overall noise impact. This process is especially useful for building a better and more efficient reward based algorithm for future learning. The base reward is a fixed incentive to reach the terminating altitude, telling the agent that completing the episode is a desired behavior. The reward functions incorporate the following factors:

1. Energy-optimization Rewards
 - a. Intermediate Reward
 - i. To incentivize the agent to progress towards its destination, the distance flown toward the destination (Δd_{OD}) during each time step is multiplied by a positive weight (w_1).
 - ii. Fe^3 models the consumption of energy for each time step. To encourage energy-efficient maneuvers, energy consumption (ΔE_t) is multiplied by a negative weight (w_2) to avoid maneuvers or paths that have high energy costs.
 - b. Termination Reward
 - i. For the energy-optimizing termination reward function, a base reward of 0.75 is added to indicate that reaching the episode's end condition is a desirable behavior
 - ii. Multiplying total energy consumed by a negative weight (w_5) creates a total energy consumption penalty, which is applied to the base reward, eliminating flight trajectories with inefficient total energy costs.
2. Noise-optimization Rewards
 - a. Intermediate Reward
 - i. Similar to the energy optimization function, the agent must progress to the destination in the noise-optimization function, so the distance flown toward the destination $\Delta(d_{OD})$ for each time step is multiplied by a positive weight (w_3)
 - ii. For noise optimization, Fe^3 will use the NPD table and calculate the SEL value (SEL_t) for each time step and multiply it by negative weight (w_4), penalizing higher noise levels.

b. Termination Reward

i. For the noise-optimizing termination reward function, a base reward of 0.75 is added to indicate that reaching the episode’s end condition is a desirable behavior

ii. The peak noise penalty takes the greatest SEL from the episode and imposes a fixed punishment to discourage situations with extreme maximum noise levels

Considering the above components, the intermediate (R_{t+1}) and termination (R_t) signals for energy optimization can be formulated as in Equation 26 and Equation 27, respectively, with the initial network weights in Table 3. Similarly, the intermediate (R_{t+1}) and termination (R_t) signals for noise optimization can be formulated as in Equation 28 and Equation 29, respectively, with the initial network weights in Table 3.

$$R_{t+1}^{\text{energy}} = r_{\text{base}} + w_1 \cdot \Delta d_{OD} + w_2 \cdot \Delta E_t \tag{26}$$

$$R_t^{\text{energy}} = 0.75 + w_5 \cdot \Delta E_{\text{total}} \tag{27}$$

$$R_{t+1}^{\text{noise}} = r_{\text{base}} + w_3 \cdot \Delta d_{OD} + w_4 \cdot SEL_t \tag{28}$$

$$R_t^{\text{noise}} = 0.75 + \text{PeakNoisePenalty} \tag{29}$$

Table 3.—Reward Function Weights

w_1	w_2	w_3	w_4	w_5
0.001	-0.002	0.001	-0.5	-0.001

8 Results

In this section, we present the relationship between climb angle, energy consumption, and acoustic impact. To illustrate these tradeoffs, Figures 8 and 9 plot the total energy consumed and the estimated Sound Exposure Level (SEL) as functions of flight path angle. Together, these figures highlight the fundamental trends that guide our analysis of noise–energy tradeoffs during the climb phase.

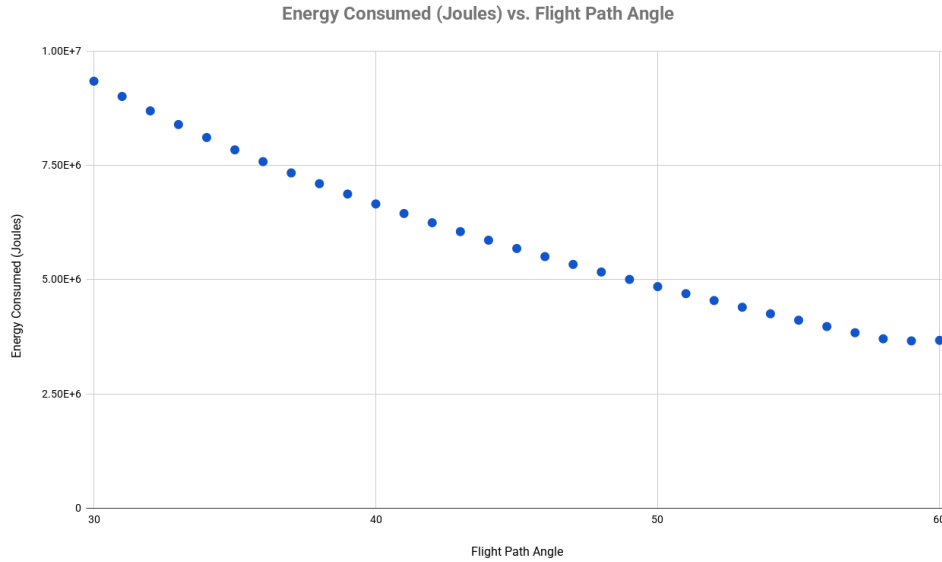


Figure 8.—Energy Consumed vs. Flight Path Angle (Steady State Climb)

The results show a clear pattern in how energy usage and noise change during a steady-state climb at different pitch angles. Figure 8 displays the energy consumed in MJ as a function of flight path angle. At lower angles (30-40°), the eVTOL consumes the most energy. This is likely because the eVTOL needs to stay in the air longer to cover the same vertical distance, flying more horizontally and using more power to maintain lift and speed. Energy consumption drops and stays relatively low from about 50° to 60°, where the eVTOL climbs efficiently without extreme power demands.

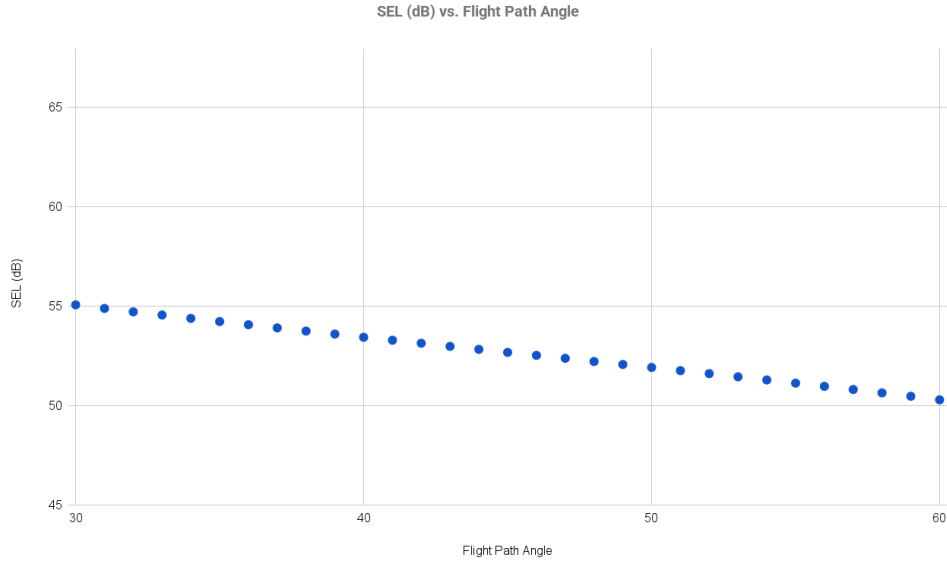


Figure 9.—Estimated SEL vs. Flight Path Angle (Steady State Climb)

A similar pattern is seen in Figure 9, which estimates the SEL produced during the climb. SEL values are highest at shallower angles, above 55 dB at 30° and steadily decrease as the pitch increases. This suggests that flying at steeper angles produces less acoustic footprint, possibly because the eVTOL spends less time flying horizontally over a given area.

Together, these trends show there’s a clear tradeoff: shallow climbs use a lot of energy and produce more noise, while steeper climbs are quieter and more efficient up to a point, but become power-hungry as the angle nears vertical. From an optimization standpoint, steeper climb angles (between 45–60°) offer the best noise reduction.

9 Conclusions

As eVTOL aircraft become more widespread with the development of UAM, noise pollution has emerged as a significant component in vehicle implementation and operational planning. This report outlines a conceptual framework using two complementary methodologies, trajectory solving through PSOPT and a deep reinforcement learning (DRL) framework, to explore different maneuvers that mitigate noise during the eVTOL climb phase. The PSOPT approach uses direct collocation to evaluate the effects of fixed climb angles on energy use and estimated SEL, while the DRL framework proposes a learning-based agent that adapts climb trajectories based on SEL and energy feedback. While only preliminary results were obtained for the optimal control method, and the implementation of the DRL agent remains ongoing, the parallel development of both approaches may inform future research for UAM trajectory design with an emphasis on both communal impact and operational feasibility. Beyond technical feasibility, the ultimate measure of success for noise-mitigation strategies lies in their effect on surrounding communities. By reducing acoustic exposure during climb, even marginal improvements in trajectory design could lessen health impacts and improve public acceptance of UAM operations.

10 Future Work

Future efforts may focus on extending the conceptual approach beyond the climb phase to also incorporate the cruise segment of the flight. The Q-function (Equation 25) will be redefined as the sum of performance metrics over both climb and cruise phases, providing a more holistic assessment of noise and energy efficiency for a full 3.7 km trajectory, beginning at vertical takeoff and extending through the top of climb to the cruise phase. In this extended model, the Q-function will be adapted to track performance across the entire 3.7 km area of interest, allowing the agent to learn policies that account for trade-offs between climb angle and the following cruise duration. A steeper climb angle results in a shorter climb phase and a correspondingly longer cruise phase, while a shallower angle leads to a longer climb and shorter cruise. Modeling these trade-offs will allow for more accurate evaluation. Calculating the SEL over a radius of 3.7 km ensures a more comprehensive assessment of noise impact, rather than being limited to the takeoff location. By extending SEL analysis across a larger radius, the framework will more accurately simulate the influence of different flight profiles on community noise exposure.

On the optimal-control side, we'll extend the PSOPT setup to cover both the climb and a short level-cruise segment. The solver will choose when the climb ends and cruise begins. To keep energy results comparable, every trajectory will cover the same total downrange distance (e.g., 3.7 km) from lift-off to the mission end point. We'll require the solution to reach the target cruise altitude and speed, and the cruise segment will hold a zero climb angle (level flight). Over this fixed distance, we'll run two separate cases: one that minimizes total energy used and one that minimizes community noise, with noise evaluated over a grid of ground receivers within 3.7 km. The DRL setup will follow the same rule (episodes end once the fixed distance is reached), and we may also report results per unit distance (e.g., MJ/km) to allow direct, like-for-like comparison.

References

1. Doo, Johnny T. and others: NASA Electric Vertical Takeoff and Landing (eVTOL) Aircraft Technology for Public Services – A White Paper. https://ntrs.nasa.gov/api/citations/20205000636/downloads/2021-08-20-eVTOL-White-Paper-Final_V48.pdf, 2021. Accessed: 2025-07-5.
2. NASA Aeronautics Research Institute: Urban Air Mobility (UAM) Concept of Operations 2.0. Version 2.0, NASA, 2020. URL https://nari.arc.nasa.gov/sites/default/files/attachments/UAM_ConOps_v2.0.pdf, accessed: 2025-07-29.
3. Whiting, T.: NASA’s Advanced Air Mobility Mission Researches Noise. <https://www.nasa.gov/centers-and-facilities/armstrong/nasas-advanced-air-mobility-mission-researches-noise/>, 2022. Accessed: 2025-07-5.
4. Daigle, M.; and Kulkarni, C. S.: Electrochemistry-based battery modeling for prognostics. *Annual conference of the PHM society*, vol. 5, 2013.
5. Li, J.; Ng, H. K.; Zheng, Y.; and Gutierreznolasco, S.: Noise exposure maps for urban air mobility. *AIAA Aviation 2021 Forum*, 2021, p. 3203.
6. Ng, H. K.: Noise impact analysis for urban air mobility in dallas-fort worth metroplex. *AIAA AVIATION 2022 Forum*, 2022, p. 3404.
7. Li, J.; Zheng, Y.; Rafaelof, M.; Ng, H. K.; and Rizzi, S. A.: AIRNOISEUAM: An Urban Air Mobility Noise-Exposure Prediction Tool. *INTER-NOISE and NOISE-CON Congress and Conference Proceedings*, Institute of Noise Control Engineering, vol. 263, 2021, pp. 474–485.
8. Valente, V. T.; Greenwood, E.; and Johnson, E. N.: An experimental investigation of eVTOL flight state variance on noise. *Vertical Flight Society’s 79th Annual Forum and Technology Display, West Palm Beach, US*, 2023, pp. 16–18.
9. Li, S.; and Lee, S.: Predictions and validations of small-scale rotor noise using ucd-quietfly. *Proceedings of the American Helicopter Society Technical Conference on Aeromechanics Design for Transformative Vertical Flight, San Jose, CA*, 2022, pp. 25–27.
10. Schatzman, N. L.; and Holly, K.: RApid Blade and Blade-Vortex InTeraction (RABBIT) Reimagined: Algorithm Improvement and User-Friendly Interface Implementation. *2024 Transformative Vertical Flight*, 2024.
11. Rizzi, S. A.; Huff, D. L.; Boyd, D. D.; Bent, P.; Henderson, B. S.; Pascioni, K. A.; Sargent, D. C.; Josephson, D. L.; Marsan, M.; He, H. B.; et al.: Urban air mobility noise: Current practice, gaps, and recommendations. , 2020.
12. Pradeep, P.; et al.: Parametric Study of State of Charge for an Electric Aircraft in Urban Air Mobility. https://ntrs.nasa.gov/api/citations/20210017230/downloads/20210017230_Pradeep_Aviation2021_manuscript_final.pdf, 2021. Accessed: 2025-07-2.
13. Laura, B.; Ronja, K.; Jochen, F.; and Eike, S.: Study On Public Acceptance of EVTOL: Safety & Noise.

14. Basner, M.; Clark, C.; Hansell, A.; Hileman, J. I.; Janssen, S.; Shepherd, K.; and Sparrow, V.: Aviation noise impacts: state of the science. *Noise and health*, vol. 19, no. 87, 2017, pp. 41–50.
15. Silva, C.; Johnson, W. R.; Solis, E.; Patterson, M. D.; and Antcliff, K. R.: VTOL urban air mobility concept vehicles for technology development. *2018 Aviation Technology, Integration, and Operations Conference*, 2018, p. 3847.
16. Patterson, M.: An Overview of NASA’s AAM Mission and Select Partnerships. *AIAA Aviation Forum and Exposition*, 2023.
17. Weather & Climate: Average Monthly Wind Speed in San Francisco. <https://weather-and-climate.com/average-monthly-Wind-speed,San-Francisco,United-States-of-America>. Accessed July 22, 2025.
18. San Francisco Metropolitan Area Weather Guide: Weather in the San Francisco Bay. <https://www.sfmex.org/wp-content/uploads/2017/03/Weather-in-the-San-Francisco-Bay.pdf>, 2017. Accessed July 22, 2025.
19. Kulkarni, C. S.; Pradeep, P.; and Chatterji, G. B.: Simulation Studies for an Urban Air Mobility Aircraft using Hardware-In-Loop Experiments. *AIAA AVIATION 2022 Forum*, 2022, p. 3648.
20. Kelly, M.: An introduction to trajectory optimization: How to do your own direct collocation. *SIAM review*, vol. 59, no. 4, 2017, pp. 849–904.
21. Pradeep, P.; Lauderdale, T. A.; Chatterji, G. B.; Sheth, K.; Lai, C. F.; Sridhar, B.; Edholm, K.-M.; and Erzberger, H.: Wind-optimal trajectories for multirotor eVTOL aircraft on UAM missions. *Aiaa Aviation 2020 Forum*, 2020, p. 3271.
22. SVANTEK: Sound Exposure Level SEL — LAE — Application — Formula. https://svantek.com/academy/sound-exposure-level/#elementor-toc__heading-anchor-2, 2025. Accessed: 2025-07-27.
23. Xue, Min: Fe3 Traffic Simulator Overview NASA Ames Technology Transfer IT Software Webinar. https://ntrs.nasa.gov/api/citations/20220012698/downloads/20220012698_Xue_AmesT2IT-Software-Fe3_presentation_final.pdf, 2022. Accessed: 2025-07-16.
24. Nickolls, J.; Buck, I.; Garland, M.; and Skadron, K.: Scalable Parallel Programming with CUDA: Is CUDA the parallel programming model that application developers have been waiting for? *Queue*, vol. 6, no. 2, Mar. 2008, p. 40–53. URL <https://doi.org/10.1145/1365490.1365500>.
25. Stevens, W. R.; and Narten, T.: Unix network programming. *SIGCOMM Comput. Commun. Rev.*, vol. 20, no. 2, Apr. 1990, p. 8–9. URL <https://doi.org/10.1145/378570.378600>.
26. Kelleher, J.: *Deep Learning*. MIT Press, 2019.
27. Póczos, B.: Introduction to Machine Learning, Reinforcement Learning. <https://www.cs.cmu.edu/~mgormley/courses/10601-s17/slides/lecture26-ri.pdf>, 2017. Accessed: 2025-07-24.

28. Oppermann, A.: Double Deep Q-Learning. <https://builtin.com/artificial-intelligence/double-deep-q-learning>, 2025. Accessed: 2025-07-27.
29. Yarramreddy, G. S.; de Alvear Cárdenas, J. I.; Pradeep, P.; Xue, M.; Lee, S.; and Kuo, V. H.: Simulation Framework for Tactical Separation Assurance Service using Deep Reinforcement Learning. <https://ntrs.nasa.gov/api/citations/20250002761/downloads/NASA-TM-20250002761.pdf>, 2025. Accessed: 2025-06-25.
30. National Oceanic and Atmospheric Administration: Data Tools: Find a Station. <https://www.ncdc.noaa.gov/cdo-web/datatools/findstation>, 2025. Accessed: July 17, 2025.

