Quadrotor-boundary and quadrotor-quadrotor interactions: analysis through reduced-order simulations

A Thesis

presented to the faculty of University of Virginia

in partial fulfillment of the requirements for the degree of

Master of Science

Lauren Bouchard

August 2021

APPROVAL SHEET

This

is submitted in partial fulfillment of the requirements
for the degree of

Author:

Advisor:

Advisor:

Committee Member:

Committee Member:

Committee Member:

Committee Member:

Committee Member:

Committee Member:

Accepted for the School of Engineering and Applied Science:

Craig H. Benson, School of Engineering and Applied Science

Abstract

Quadrotors, a class of unmanned aerial vehicles (UAVs) or drones, are gaining popularity across a variety of use cases. Many proposed applications require multiple UAVs to fly in close proximity to obstacles like surfaces or other drones, where the differences in the flowfield can disrupt the UAV's flight path. This work investigates two near-obstacle flight considerations.

In the first, a quadrotor hovers near a ground or ceiling boundary. While the increased lift in those regions may result in energy savings, the proximity to the boundary presents a risk of crashing. I explore this tradeoff, as well as the importance of the accuracy of the boundary model, via reduced-order simulations. Simulations with example parameters suggest that quadrotors can safely fly roughly one half rotor-radius closer to the ground than we would expect without having considered the ground effect. Safe near-ceiling flight should be about one rotor-radius further from the boundary than would otherwise be considered safe. This information can provide guidance for future UAV system designers.

In the second near-obstacle flight scenario, a quadrotor flies in close proximity to another flying quadrotor. This is relevant in swarming applications, where multiple UAVs coordinate their actions, often in close proximity. I investigated one such situation: a quadrotor flies horizontally closely above or below a hovering quadrotor, and its path is altered as a result. This was tested via simulations that combine potential flow solutions with a simple quadrotor dynamics model. These simulations were compared to the results of an analogous experiment conducted by Esen Yel. A variety of flow models and fit levels were assessed. The heuristic (unfitted) simulations almost always over-predicted the moving quadrotor's deflection. There was variety in performance across flow models: the flights through flow fields informed by particle image velocimetry performed more realistically than the other models. Iterative fitting improved accuracy to a consistent level across all models.

Together, these studies highlight the importance of local airflow on UAV path-planning. The models presented here are simple physics-based simulation methods that do not require extensive experimentation or computational resources. Future efforts might expand upon this in several directions; work is already underway at UVA MAE to characterize the impact of a rotor's blade twist and size on near-boundary control dynamics.

Acknowledgements

I would first like to thank my advisor, Dan Quinn. I am particularly grateful for his guidance on the design of my research projects. My work would not be nearly as coherent without the many hours he has generously spent helping me decide which simulations to write and how to explain them. I would also like to thank Darius Carter and Esen Yel for supplying the experimental data that give my work important context and for dutifully answering my many questions about their experimental setups. Thank you to the Smart Fluid Systems Lab and my friends in the Link Lab for their technical advice and support as well as their company throughout my time as a graduate student. Last, I'd like to thank Chloe Dedic and Chris Goyne for sitting on my committee.

My work was partially funded though the National Science Foundation NRT program under Grant No.1829004.

Contents

Abstract ii
Acknowledgements iii

1 Introduction 1
1.1 Background: Quadrotors and their uses 1
1.2 Motivation: Near-obstacle quadrotor flight 2
1.3 Motivation: Reduced-order simulations 3
1.4 Scope of work 4

2 Hovering Near a Boundary 5
2.1 Boundary Lift Models 5
2.2 Numeric Simulation: 1D rigid body dynamics with boundary effect 6
2.3 Crash Propensity Findings 7
2.4 Efficiency Tradeoffs 9

3 Flight through another quadrotor's wake 10
3.1 Experimental design 10
3.2 Numeric Simulation: 2D rigid body dynamics with potential flow 11
3.3 Experimental findings and modeling implications 16
3.4 Simulation Results 17

4 Conclusion 21
4.1 Key Findings 21
4.2 Contributions 22
4.3 Future Work 22

References 23

1 Introduction

1.1 Background: Quadrotors and their uses

Unmanned aerial vehicles (UAVs) have existed for over a century. Historic UAVs were typically fixed-wing aircraft used for offensive warfare. Many newer models employ rotor-based designs as their high maneuverability allows for use in a variety of close-range situations, notably increasing their civilian utility [1]. Quadrotor (four-rotor) configurations can be seen today in applications ranging from aerial photography [1] to coordinated light shows [2].

Quadrotors come in an assortment of sizes (Figure 1.1). The DJI Mini 2, a popular consumer photography drone, is 213 mm across and has a rotor diameter of 119mm. Professional models are often larger - the same manufacturer's professional photography UAV, the Inspire 2, is 605 mm diagonally with a rotor diameter of 380 mm. Other applications benefit from smaller drones: the Crazyflie, which is marketed for research and education, is considered a "micro aerial vehicle" at only 130 mm across with 46 mm diameter rotors. The smaller size means crashes are less impactful, giving researchers flexibility to test riskier flight maneuvers and control algorithms.

Figure 1.1: **Left: The DJI Inspire 2, a professional photography quadrotor. Right: The Crazyflie Micro Aerial Vehicle** *Left image from Elmekkaoui abdelghani via Wikimedia (CC BY-SA 4.0).*

Across all sizes, the development of accurate and safe autonomous UAV control is necessary for the next generation of drone applications. While aerial photographers often command their drones manually via handheld controllers, preplanned flight control systems do exist and are used in, for example, drone light shows. Notably, light shows allow for a large clearance between the individual quadrotors and between the quadrotors and obstacles, meaning the drone can stray far from its path without significant consequences. For quadrotors to fly in crowded spaces where flight close to obstacles is inevitable, they will have to precisely adhere to their prescribed trajectories in order to avoid collisions.

1.2 Motivation: Near-obstacle quadrotor flight

Many upcoming applications require close-proximity flight that would demand a small margin of error in each vehicle's flight path. This especially applies to flight in crowded urban environments, for example package delivery [1]. At the same time, flights close to surfaces or other drones adhere less to their prescribed trajectories because airflow is impacted by the obstacles. To ensure safety in close proximity flight, future applications will require an understanding of the impact of a nearby surface (the "boundary effect") or another UAV on the trajectory of the vehicle.

Flight near boundaries is practically unavoidable, especially in congested urban environments. First, UAVs generally take off from or land on a ground surface. While autonomous quadrotors usually have sensors to facilitate landing, an understanding of the local fluid flow could allow for sensorless landing [3, 4, 5]. Some applications, such as bridge inspection, require a quadrotor to approach a ceiling surface [6]. In other situations, a quadrotor might be able to fly a more direct path if it does not have travel far around obstacles such as tables or rooftops. Aside from a shorter path distance, the flight could also be more efficient as a result of the altered lift in boundary regions.

When a quadrotor flies close to a ground or ceiling, it experiences an increase in lift compared to that of the same rotor speed far from the boundary. For that reason, flight requires less energy in these regions [7, 8], which is important because many UAV applications are constrained by the vehicle's battery capacity. The increase in lift has stability implications, too: the ceiling acts as an attractor, pulling the quadrotor into a positive feedback loop, while the ground has a cushioning effect since closer flight pushes the quadrotor upward at a larger magnitude [9, 10].

Other applications require coordinated flight of several UAVs, often referred to as "swarming"[11]. This has been proposed for use in a wide variety of applications including cooperative search [12, 13] and urban traffic monitoring [14, 15]. Unlike current swarms (eg. a light show), many upcoming applications require the drones to fly in close proximity to one another, where the fluid disturbance from one vehicle can have a significant impact on the flight path of another.

Not much is known about quadrotor-quadrotor flow interactions. Seeing as each quadrotor forces air downward to produce lift, it is logical to expect a UAV flying above or below a hovering vehicle might experience a downward deflection in its path. Important factors may include the speed of the moving quadrotor and the relative altitude between the two. Additionally, the deflection might look different when flying above the quadrotor, where the flow is laminar, compared to below the quadrotor, where flow is concentrated into a turbulent downwash. While a model to predict this deflection is vital to ensuring safe flight, researchers have yet to characterize the phenomenon.

1.3 Motivation: Reduced-order simulations

While all models are incomplete representations of reality, a reduced-order model intentionally sacrifices some accuracy for the benefit of simplicity. One example in the context of this work is flow approximation. Many studies estimate flow with computational fluid dynamics (CFD) methods that solve the Navier-Stokes equations at a grid of points. These methods can be quite accurate, but they require significant computational resources. Instead, the reduced-order flow models in Chapter 3 use potential flow solutions (for example, a unidirectional flow) as a rough approximation of a flowfield.

There are drawbacks to reduced-order models. Most notably, an oversimplification might omit important physics and, as a result, the model outputs could be too inaccurate for their purpose. In the above example, a potential flow simplification assumes that the flow is inviscid and irrotational, meaning the model will not capture any turbulence. A CFD method, on the other hand, can predict turbulence. Additionally, reduced-order models often require some degree of training or tuning with real data in order to partially compensate for missing physics in the model.

Despite these drawbacks, simplified models are helpful in many situations, especially those that benefit from small computational requirements. For instance, a quadrotor's onboard processor could execute a reduced-order simulation to inform its flight path in real time. Fast simulations are also necessary for purposes where each simulation needs to be iterated multiple times, as I do in Chapter 2 to assess the impact of random forcing on a quadrotor's flight.

Simple simulations can also give researchers an opportunity to investigate the implications of existing analytical or experimental models in the context of larger systems. In Chapter 2, I add boundary models to a reduced-order simulation of rigid-body dynamics and proportional-integral-derivative control. This method facilitates the study of an important outcome, crash propensity, based on the whole system's behavior given a specific boundary model. In this situation, a detailed flow solution does not make sense as the question of interest is the system's sensitivity to the accuracy of the flow model itself.

Additionally, while model errors aren't generally seen as desirable, the shortcomings of a reduced-order model can actually reveal important information. A high error usually implies that the physics programmed into the model is incorrect or omits important details. A researcher could leverage this to verify and refine their understanding of a system. In Chapter 3 I use model error to make inferences about the flowfield in close proximity quadrotor-quadrotor flight.

Because of their simplicity, reduced-order models can also be used as the basis for guidelines for drone operators. While it would be unreasonable to expect a typical UAV pilot to run a high-fidelity CFD simulation, it might be feasible for the pilot to run a simple physics-based model to, for example, heuristically estimate how much clearance a quadrotor needs when flying below another quadrotor. Ultimately, if the reduced-order model is highly accurate, its grounding in physics paves the way for analysis to prove safety of a system.

1.4 Scope of work

Here, I investigate two specific near-obstacle quadrotor flight scenarios. For each, I present a reduced-order model designed to meet the objectives of the study. I use experimental data to inform my simulations and, in Chapter 3, assess the accuracy of the simulations. I place emphasis on analyzing each simulation in context of its possible use cases.

Chapter 2 covers an investigation of the impact of the boundary effect on a quadrotor's crash propensity. The one-dimensional simulation predicts a quadrotor's altitude and incorporates a ground model refined by Darius Carter, who is also a graduate student at UVA MAE. The purpose of the simulation is to investigate the importance of a precise boundary model on a quadrotor's crash propensity and energy efficiency.

Chapter 3 outlines a two-dimensional simulation of quadrotor flight through another quadrotor's wake, which is approximated by canonical potential flow solutions. The simulated quadrotor paths are compared to experimental data provided by Esen Yel and Nicola Bezzo of UVA Systems Engineering. In addition to the flow models, I also compare a variety of fitting techniques, starting with heuristically-determined model parameters. This analysis informs the potential use of a combined potential flow and rigid body dynamics simulation technique.

2 Hovering Near a Boundary

2.1 Boundary Lift Models

Rotorcrafts experience increased lift near a boundary. This was first studied in helicopters, where researchers analytically modeled the ground by superimposing a flow source at the rotor's center with a copy that has been "mirrored" an equal distance below the ground plane (the method of image) [16]. This resulted in a model for near-ground lift L compared to lift far from the boundary L_∞ as a function of altitude z scaled by rotor radius r, $\hat{z} \equiv z/r$:

$$\frac{L}{L_\infty} = \frac{1}{1 - \frac{1}{16\hat{z}^2}}. \qquad \text{(Eq. 2.1)}$$

Recent efforts have extended this theory to accommodate four rotors. Sanchez-Cuevas et al [7] modeled four sources on a quadrotor with a rotor-rotor distance ℓ scaled by rotor radius, $\hat{\ell} \equiv \ell/r$:

$$\frac{L}{L_\infty} = \frac{1}{1 - \frac{1}{16\hat{z}^2} - \frac{\hat{z}}{\sqrt{\left(\hat{\ell}^2+4\hat{z}^2\right)^3}} - \frac{\hat{z}}{2\sqrt{\left(2\hat{\ell}^2+4\hat{z}^2\right)^3}} - \frac{\hat{z}}{2K_b\sqrt{\left(\hat{\ell}^2+4\hat{z}^2\right)^3}}}. \qquad \text{(Eq. 2.2)}$$

The last term uses a semi-empirical coefficient K_b to account for an interaction between the individual rotors: the "fountain effect" is an upward flow located beneath the center of the quadrotor. Sanchez-Cuevas et al found that the model best agreed with experimental data when $K_b \approx 2$ [7].

This model depends on several assumptions. To estimate lift, it assumes that the downwash is constant across the rotor, that the rotor is thin, and that the distance from the ground is greater than $\hat{z} \approx 0.25$, where there is a singularity in the function. The use of a flow source comes from potential flow theory, meaning all its assumptions also apply here, including that the flow is considered inviscid and incompressible.

The models based on these assumptions proved reasonable when applied to full-sized quadrotors [7]. Because micro quadrotors are much smaller than the helicopters that inspired the theory, however, viscous effects might no longer be negligible. This prompted further experimental analysis. Carter et al conducted an experiment with a Crazyflie MAV ($r = 23\,mm$) which confirmed that the model (Equation 2.2), with the same empirical coefficient $K_b = 2$, holds well until $\hat{z} \approx 0.5$ [9].

Rotorcrafts also experience lift near a ceiling. This is a relatively new topic of research and has not been investigated as extensively as the ground effect. An analytical approximation [8] uses momentum theory on the volume immediately surrounding a quadrotor:

$$\frac{L}{L_\infty} = \frac{1}{2} + \frac{1}{2}\sqrt{1 + \frac{1}{8\hat{z}^2}} \quad \text{(Eq. 2.3)}$$

Results from experiments with small rotors ($r = 23\,mm$ [9, 8] and $r = 50\,mm$ [8]) are in good agreement with the model at distances further than half a rotor radius from the ceiling. Closer distances resulted in a smaller increases in lift than the values predicted by the model [9].

This limited discrepancy between the model predictions and experimental data raises an important question: what is a sufficient level of model accuracy? While the answer to this questions depends on the specific UAV application, this chapter focuses on the implications of the model and its accuracy on vehicle crash propensity, a common consideration between most use cases. Because future flights may wish to leverage the boundary effect to conserve energy, I also compare a given model's predicted crash likelihood to the relative energy savings of near-boundary flight.

2.2 Numeric Simulation: 1D rigid body dynamics with boundary effect

To investigate the implications of the boundary effect on UAV safety, I created a reduced-order simulation of near-boundary quadrotor dynamics to determine relative crash propensities. The simulation compares crash rates of different boundary models across a range of flight altitudes, using a Crazyflie's dimensions ($r = 23$ mm, $\ell = 90$ mm) and weight (27 g) as an example. I also analytically estimated the relative power required to maintain a given altitude. This information, together with the crash propensities, highlights efficiency-safety tradeoffs inherent in near-boundary flight.

The reduced-order model considers three forces acting on the Crazyflie: the lift force (L), the vehicle's own weight (mg), and random vertical forces chosen to simulate disturbances (L'). These random forces were generated by a white Gaussian noise function (mean of 0 N and standard deviation of 0.015 N). From a force balance, the vehicle's acceleration follows as

$$\ddot{z}(t) = \frac{L(z(t)) + L'(t) - mg}{m}, \quad \text{(Eq. 2.4)}$$

that is solved using a Euler integration (time step of 0.1 s) to determine velocity and position over 120 s. At every time step, L was recalculated as a function of the vehicle's current altitude z and the rotor frequency f (the control variable).

To estimate L at distances far from a boundary, the simulations used blade element theory: $L = 4C_L\rho f^2 r^4$, where C_L is the lift coefficient of the rotor blade. The 4 is to account for the four rotors, and I held C_L constant at 1.6 (in doing so, I assumed lift coefficient to be independent of Reynolds Number). This lift force served as L_∞ for the calculation of the near-boundary lift indicated by the Hsaio-Chirarattananon and Sanchez-Cuevas models (Eqs. 2.2 and 2.3). I used Sanchez-Cuevas's empirically fit $K_b = 2$ in the ground effect model [7]. To avoid unreasonably large lift values (and the singularity near $\hat{z} = 0.25$ in the ground model), I bounded L/L_∞ to its value computed at $\hat{z} = 0.5$. To assess the sensitivity of vehicle safety to the accuracy of the boundary model, I also simulated vehicle dynamics with boundary models scaled by powers of 0.6, 0.8, 1.2, and 1.4. For example, the lift in one simulation would be computed as $(L/L_\infty)^{1.4}$, such that the lift is always greater than L/L_∞ indicated by the model yet still approaches L_∞ at distances far from the boundary.

The vehicle was commanded to maintain a target height (measured as distance from the rotor to the boundary) throughout each simulation using a simple proportional-integral-derivative (PID) controller on the rotor frequency. At each timestep, the frequency was determined by

$$f = k_P(z - z_0) + k_D\dot{z} + k_I \int_0^t (z - z_0)\mathrm{d}t + f_0, \qquad \text{(Eq. 2.5)}$$

where z_0 is the target height, f_0 is a frequency offset, and k_P, k_D, and k_I are the controller gains. Gains were held constant across all simulations. The k_P and k_D gains were chosen such that a quadrotor hypothetically launched from the ground would quickly reach its desired altitude with minimal overshoot; I decided upon $k_P = 316\,\mathrm{m}^{-1}\mathrm{s}^{-1}$ and $k_D = 316\,\mathrm{m}^{-1}$. Because I ultimately wanted the vehicle to start at and maintain a given height throughout our simulations, I added a frequency offset sufficient to counter the Crazyflie's own weight (adjusted for increased lift due to the boundary effect in applicable simulations). While an integral gain was still necessary to account for any asymmetry in the disturbances, including this offset meant a relatively small integral gain of $31.6\,\mathrm{m}^{-1}\mathrm{s}^{-2}$ was sufficient to consistently maintain a desired average altitude. The minimum rotor frequency was also bounded at 0, i.e. the rotors could not spin backwards to provide downward thrust in the event the altitude far exceeded the target height.

Vehicle target heights ranged from $\hat{z}_0 \approx$ 0.5 to 2 above ground and $\hat{z}_0 \approx$ 1 to 2.5 below the ceiling (where $\hat{z}_0 \equiv z_0/r$). If at any time during the simulation $\hat{z} \leq 0$, I considered the vehicle to have crashed into the boundary and thus stopped the simulation, adding to a counter whenever this happened. Simulations of each combination of setpoint and boundary model magnitude were repeated 1000 times to determine an average crash rate at each parameter combination.

2.3 Crash Propensity Findings

My reduced-order simulations demonstrate that the ground has a stabilizing effect while the ceiling has a destabilizing effect. Random fluctuations cause the quadrotor to deviate from its target height, which could potentially lead to crashes with a nearby boundary. However, as a quadrotor approaches the ground, the heightened lift pushes the quadrotor upwards and prevents a crash (Fig. 2.1*a*). In contrast, approaching a ceiling leads to higher forces toward the boundary, which can result in a crash (Fig. 2.1*b*). To explore the likelihood of crashes, I aggregated hundreds of trials and looked at average crash rates.

On average, the simulated quadrotor crashes less near the ground when a ground model is included in the simulation. With no boundary model, random fluctuations cause crashes as high as $\hat{z}_0 \approx 1.5$ (Fig.

2.1*c*). When a ground model is added, the quadrotor can be about half of a rotor radius closer to the ground before this rise in crash rate. Scaling the ground model causes only slight changes in the $\hat{z}_0$ value at which this rise occurs.

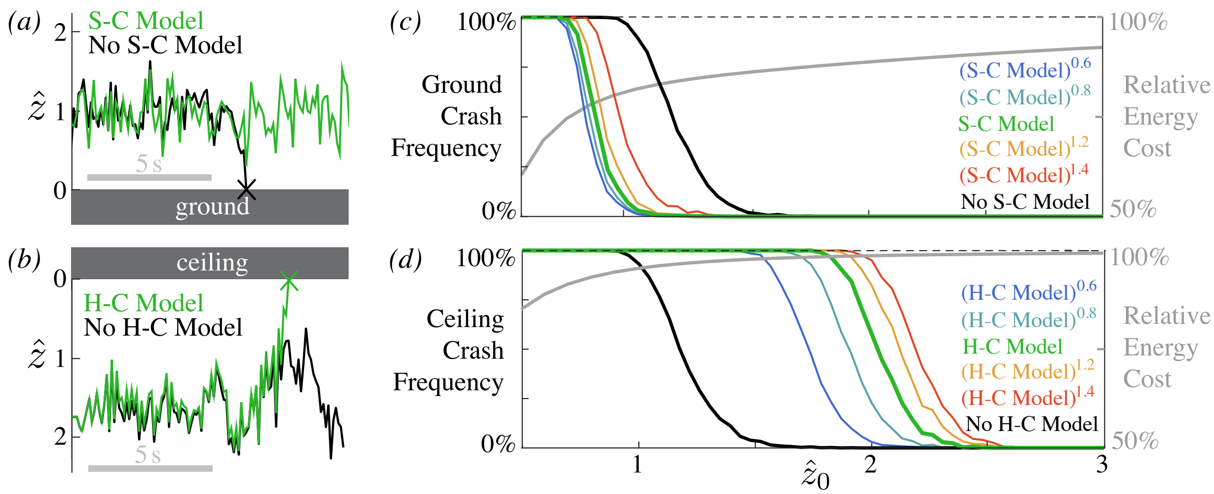

Figure 2.1: Reduced-order quadrotor simulation. a,b) Sample time histories with and without the ground/ceiling models. c,d) Average crash frequencies with and without ground/ceiling models (scaled by exponentiation).

Unlike a quadrotor near the ground, a quadrotor near the ceiling experiences more crashes. On average, the quadrotor is likely to crash into the ceiling when $\hat{z}_0 <\approx 2$ (Fig. 2.1*d*). Scaling the models has a stronger effect on the safe $\hat{z}_0$ range near the ceiling than it does for the safe range near the ground. The effects differ in magnitude because they are caused by different mechanisms. The ground acts as a cushion that pushes the quadrotor away; crashes require large random fluctuations. The ceiling acts as an attractor, pulling the quadrotor into a positive lift feedback loop; crashes are inevitable unless the controller can reverse course in time.

Note that the value of the safe/unsafe $\hat{z}$ values depends on the disturbance intensity (L') and the PID gains injected into our model. Varying the intensity or the gains would rescale the $\hat{z}$ values in Fig. 2.1, though the relative positioning of the curves – and therefore our conclusions about crash frequency – would be unaffected. If the equation of motion (Eq. 2.4) were non-dimensionalized, say by using rotor radius and gravity to scale lengths and times ($\hat{z} \equiv z/r$, $\hat{t} \equiv t/\sqrt{r/g}$), it could be written as

$$\ddot{\hat{z}} = \left(\alpha(\hat{z} - \hat{z}_0) + \beta\dot{\hat{z}} + \gamma \int_0^{\hat{t}} (\hat{z} - \hat{z}_0)\mathrm{d}\hat{t} + \delta\right)^2 + \frac{L'}{mg} - 1, \qquad \text{(Eq. 2.6)}$$

where the four dimensionless groups α, β, γ, and δ are the controller gains and offset scaled by physical variables: $k_\mathrm{P}\sqrt{4C_\mathrm{L}\rho r^6/(mg)}$, $k_\mathrm{D}\sqrt{4C_\mathrm{L}\rho r^5/m}$, $k_\mathrm{I}\sqrt{4C_\mathrm{L}\rho r^7/(mg^2)}$, and $f_0\sqrt{4C_\mathrm{L}\rho r^4/(mg)}$, respectively. If lift coefficient were not assumed to be independent of the blade-tip Reynolds Number (see Numerical Methods), Reynolds number would be an additional dimensionless group that includes viscosity ($\rho f r/\mu$, where μ is viscosity). Therefore, while our simulations were run with variables specific to a Crazyflie, they could be tested for quadrotors more generally by tuning α, β, γ, δ, and the disturbance function scaled by vehicle weight (L'/mg).

2.4 Efficiency Tradeoffs

In contrast to crash propensity, the relative energy cost of hovering near a boundary can be determined analytically since it reduces to a function of our boundary model. I combined the lift models with blade element theory, which provides the mechanical power generated by a rotor as $C_P \rho f^3 r^5$, where C_P is the power coefficient. For the small advance ratios ($J \equiv \dot{z}/(2rf)$) of hovering, C_L and C_P are relatively constant (e.g. < 5% change for $J < 0.2$ for a typical rotor McCormick1995). The ratio of power consumed in two different flight conditions 1 and 2 is therefore $\approx (f_1/f_2)^3$. Inverting $L = C_L \rho f^2 r^4$ gives the frequency f required for hover as a function of $\sqrt{L}$. Therefore, the mechanical power generated near a boundary compared to the power far from the boundary ("relative energy cost") is $(L/L_\infty)^{-3/2}$, where L/L_∞ follows from Eqs. 2.2 and 2.3.

For comparison, I also plotted the relative energy costs of near-boundary flight in order to highlight the tradeoff between safety and efficiency. This relation makes it clear why near-boundary flight is more efficient: as $\hat{z}$ drops, L/L_∞ goes up, requiring a smaller rotor frequency and less energy to maintain altitude. Chances of crashing increase with smaller $\hat{z}_0$, however, so accurate models are critical for balancing safety and efficiency near the ground/ceiling (Fig. 2.1*c,d*).

3 Flight through another quadrotor's wake

3.1 Experimental design

We currently do not have models that predict the deflection of a moving quadrotor as it approaches another vehicle. Studies to date have provided physics-based models for hovering quadrotors near boundaries such as a floor or ceiling [9, 7, 5], but little is known about the form of the flow between two neighboring vehicles. Additionally, the large number of factors influencing the dynamics of non-stationary flight preclude extensive analysis using typical experimental or numeric methods.

There are, however, tools that can approximate different parts of this complicated system. Knowledge of a quadrotor's control algorithm can be paired with rigid body dynamics and blade element theory to simulate real flight behavior [17, 18]. Separately, solutions to potential flow are often used to approximate high Reynolds number flows including those around aircraft. While turbulent flows, such as in the wake of a quadrotor, are more complicated, classic closed-form functions can still be of use for rough modeling purposes when the time scale of interest is much longer than that of the turbulence. A comparison of particle-image velocimetry (PIV) data to canonical flow models suggests that the models can roughly represent the regions of the airfield around a quadrotor (Figure 3.1).

Here, I combine a rigid body modeling technique with a simplified flow approximation by using the flow to estimate a force on the quadrotor. This force is a direct input to the dynamic equations. The goal is for this combined simulation to accurately predict the path disturbance of a quadrotor as it flies in close proximity to another aircraft.

To explore the accuracy of this method, I collaborated with Esen Yel and Nicola Bezzo of UVA Systems Engineering (and my advisor, Dan Quinn) to design an experiment. We settled on a test where a quadrotor was commanded to move horizontally above or below a hovering quadrotor at two flight velocities across eight relative altitudes. I compare the experimental results to my simulated flight paths, which employ six different flow models. My simulations first use heuristically-determined parameters based on physical aircraft dimensions and common gain tuning methods. I then explore a variety of fitting scenarios to determine how the simulation inputs relate to the accuracy of the model predictions.

My collaborators conducted the experiment in the fall of 2020. They collected experimental data using two AscTec Hummingbird quadrotors (mass of 526-530 g and rotor diameter of 20.32 cm; see Table 3.1). The experimental setup used to collect data is presented in Figure 3.2. The first quadrotor

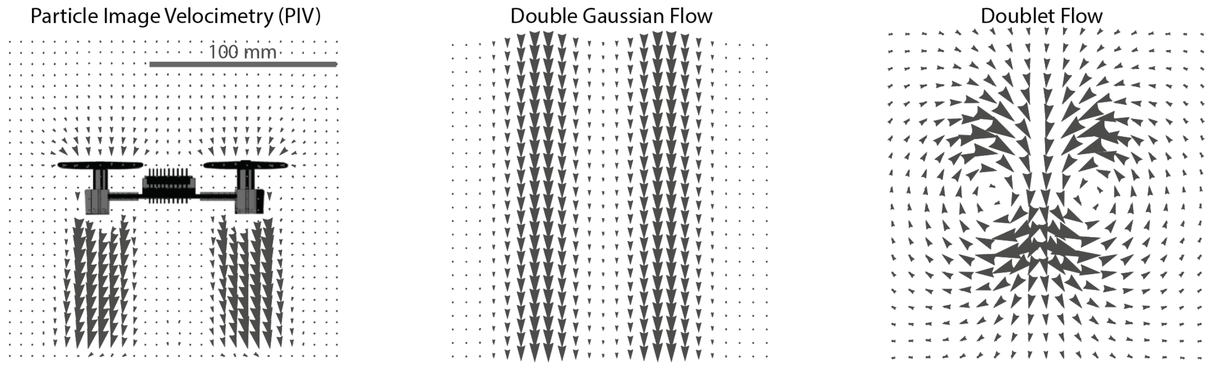

Figure 3.1: **Comparison of flow measurements and simplified flow models. Left: Particle Image Velocimetry flow measurements around a quadrotor [9]. Center: unidirectional flow with a magnitude determined by two Gaussian functions roughly approximates the region below the quadrotor. Right: a doublet flow model roughly approximates the region above each rotor.**

hovers in the middle of the room (red dot, $x = 0\,\mathrm{m}$) and the second one starts its motion from an initial position (blue dot, $x_1 = -2\,\mathrm{m}$) and moves to its goal position (green dot, $x_2 = 2\,\mathrm{m}$). The moving quadrotor follows a minimum-jerk trajectory [19, 20] with two different average commanded speeds: $\bar{\dot{x}}_c$ = 0.3 and 0.8 m/s. To analyze the effects of the airflow both below and above the hovering quadrotor, this motion was repeated with eight commanded altitude differentials between the two quadrotors: -1.5, -1.2, -1.0, -0.5, -0.3, 0.2, 0.5, and 1.0 m. (Ground altitudes between 0.5 m and 2 m; the ceiling is roughly 3 m high.) In total, data was collected from 16 different trajectories. These are visualized in Figure 3.5.

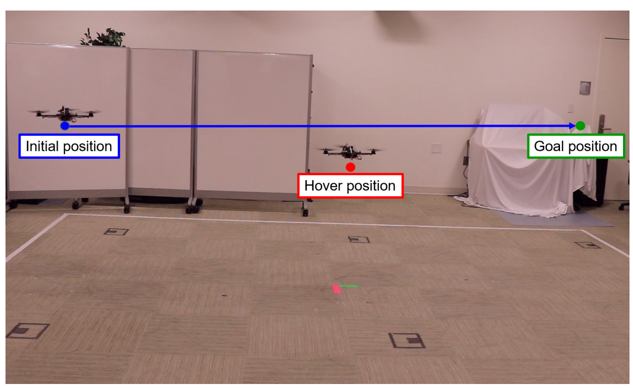

Figure 3.2: **The quadrotor attempts to fly from its initial position $\langle x_1, 0\rangle$ to a goal position $\langle x_2, 0\rangle$, potentially encountering a flow disturbance along the way. Photo courtesy of Esen Yel.**

3.2 Numeric Simulation: 2D rigid body dynamics with potential flow

To simulate a quadrotor's flight, I employed a reduced-order simulation of the vehicle's dynamics and controls. While several quadrotor control models exist [17, 19, 20], I chose to use a cascaded Proportional-

Derivative (PD) controller similar to prior work [21] because of its simplicity and tracking efficacy. An external forcing representing a flow field is included in the dynamic equations of this model.

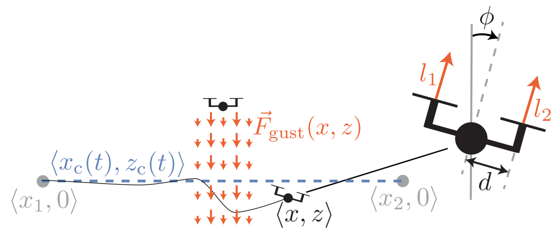

Figure 3.3: **An illustration of the quadrotor's fully-defined state as it attempts the desired trajectory** $\langle x_{\mathrm{c}}(t), z_{\mathrm{c}}(t)\rangle$

The quadrotor is modeled in two dimensions, so its state is fully determined by its position ($\langle x, z\rangle$), velocity ($\langle \dot{x}, \dot{z}\rangle$), roll angle ($\phi$), and roll velocity ($\dot{\phi}$) (Figure 3.3). The quadrotor's other relevant parameters are its mass m, rotor arm length d, roll moment of inertia I, and the net and differential lift of its rotors, L_{net} and L_{dif}, which it uses to control its position. The accelerations of the quadrotor are governed by linear and angular momentum balances:

$$\ddot{x} = \frac{L_{\text{net}}\sin\phi + F_{\text{gust},x}}{m}, \quad \text{(Eq. 3.1)}$$

$$\ddot{z} = \frac{L_{\text{net}}\cos\phi + F_{\text{gust},z} - mg}{m}, \quad \text{(Eq. 3.2)}$$

$$\ddot{\phi} = \frac{dL_{\text{dif}}}{I}, \quad \text{(Eq. 3.3)}$$

where $F_{\text{gust},x}$ and $F_{\text{gust},z}$ are external forcings from flow disturbances and mg is the vehicle's weight. I considered cases where the vehicle started at rest, i.e. $\dot{x}(0) = \dot{z}(0) = \dot{\phi}(0) = 0$.

To create the commanded flight path, $\langle x_{\mathrm{c}}(t), z_{\mathrm{c}}(t)\rangle$, I used a trajectory generator that minimizes the jerk to fly between two prescribed locations: from $\langle -2, 0\rangle$ to $\langle 2, 0\rangle$. The quadrotor uses nested PD control to attempt this trajectory. During flight, the commanded roll angle is calculated based on errors in horizontal position:

$$\phi_{\mathrm{c}} = K_{\mathrm{P},x}(x - x_{\mathrm{c}}) + K_{\mathrm{D},x}(\dot{x} - \dot{x}_{\mathrm{c}}), \quad \text{(Eq. 3.4)}$$

where $K_{\mathrm{P},x}$ and $K_{\mathrm{D},x}$ are constant proportional (P) and derivative (D) gains. Then, the average and differential lift are calculated based on errors in vertical position and roll angle:

$$L_{\text{tot}} = K_{\mathrm{P},z}(z - z_{\mathrm{c}}) + K_{\mathrm{D},z}(\dot{z} - \dot{z}_{\mathrm{c}}) + mg, \quad \text{(Eq. 3.5)}$$

$$L_{\text{dif}} = K_{\mathrm{P},\phi}(\phi - \phi_{\mathrm{c}}) + K_{\mathrm{D},\phi}\dot{\phi}, \quad \text{(Eq. 3.6)}$$

where the K's are the gains for Z and ϕ dynamics. The $+mg$ is needed because without the integral gain the PD controller would otherwise reach equilibrium below the commanded altitude.

Equations 3.1-3.6 create a coupled set of second-order differential equations for $x(t)$ and $z(t)$.

Solutions are fully determined by the choice of the flow disturbance, the quadrotor's parameters (m, I, d), the choice of gains, and the commanded trajectory. I note that these variables could be arranged into twelve dimensionless groups, suggesting that our predictions could be applicable to other parameter values - e.g. quadrotor size - as long as the dimensionless groups match those of this study. I solved the system numerically using a forward-stepping Euler method ($\Delta t = 0.025\,\mathrm{s}$). Figure 3.4 shows sample simulations of a quadrotor flying through a downward jet. As expected, the quadrotor experiences a drop in altitude.

All the simulations used physical parameters and trajectories designed to approximate those of the Hummingbird quadrotor used in our experiment and are summarized in Table 1. Esen massed the two quadrotors (526 and 530 g) and used the average of the two masses in our simulation. The arm length d was measured from the center of the rotor to the center of the vehicle. For a ballpark estimate of the quadrotor's moment of inertia I, I computed the moment of inertia if the vehicle's mass were evenly distributed across a square with a side length of 6 cm, roughly twice the side length of the central connection of the quadrotor (3 cm).

The baseline simulation used heuristically-determined gains. I modified the Ziegler-Nichols tuning method [22] to accommodate our cascaded PD controller. I first tuned the vertical (z) gains by increasing the proportional gain until the quadrotor launched to a height of 1 m with an oscillation period of 2 s. This is the ultimate proportional gain, K_u, and period, T_u, required by common tuning methods. The proportional and derivative gains were then determined from $K_{\mathrm{D},z} = 0.6K_{\mathrm{u}}$ and $K_{\mathrm{D},z} = 0.1K_{\mathrm{u}} * T_{\mathrm{u}}$ respectively. The ϕ gains were tuned similarly, using a desired angle of 15 degrees and oscillation period of 1 s. Lastly, the horizontal (x) gains were chosen as 1/10th of the vertical gains since vertical motion is far more sensitive to small changes in our quadrotor's pitch. The gains used in my baseline simulation are reported in Table 3.1.

Table 3.1: **Physical parameters and control gains used in the simulations.**

Variable	Value	Description
m	528 g	Mass
g	$9.81\,\mathrm{m/s^2}$	Gravitational acceleration
d	19 cm	Rotor arm length
I	$1.58 \times 10^{-4}\,\mathrm{kg\,m^2}$	Estimated roll moment of inertia
$K_{\mathrm{P},x}$	$0.3\,\mathrm{m^{-1}}$	Horizontal proportional gain
$K_{\mathrm{D},x}$	$0.1\,\mathrm{s\,m^{-1}}$	Horizontal derivative gain
$K_{\mathrm{P},z}$	$3\,\mathrm{kg\,s^{-2}}$	Vertical proportional gain
$K_{\mathrm{D},z}$	$1\,\mathrm{kg\,s^{-1}}$	Vertical derivative gain
$K_{\mathrm{P},\phi}$	$0.019\,\mathrm{N}$	Roll proportional gain
$K_{\mathrm{D},\phi}$	$0.0032\,\mathrm{N\,s}$	Roll derivative gain

I flew the simulated quadrotor through six different flow field shapes (Figure 3.4). These models provide horizontal and vertical velocity components (u and v, respectively) as functions of the relative position of the two quadrotors (x_d and z_d). The speeds u and v are ultimately used to estimate the force on the aircraft, $F_{\mathrm{gust},x}$ and $F_{\mathrm{gust},z}$, for use in equations 3.1 and 3.2.

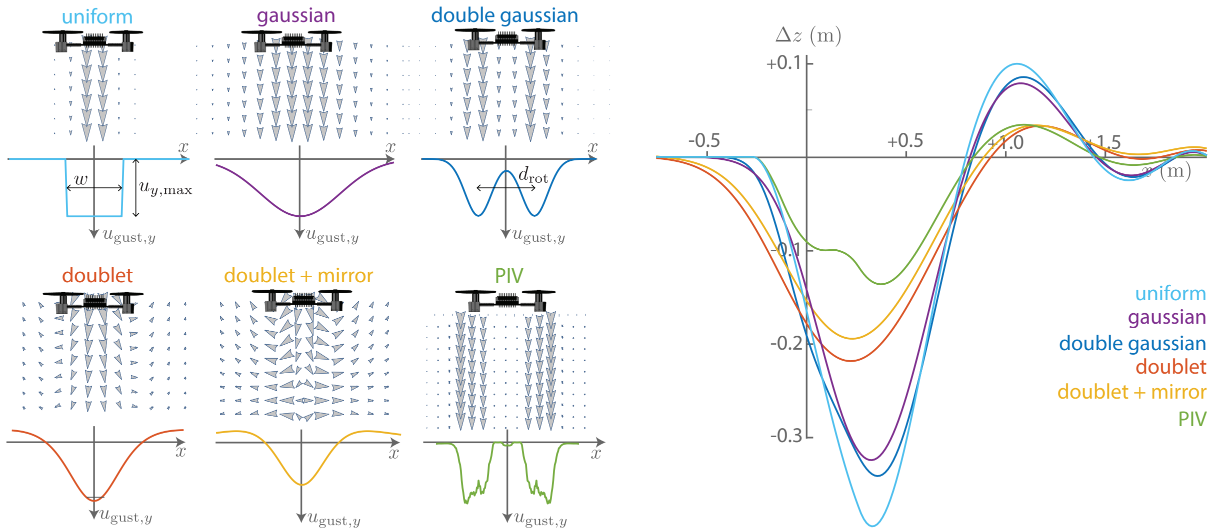

Figure 3.4: **Left: vector field representations of each of the six flow field models used for the simulation. Right: example traces of simulated quadrotor flight through each of the flow fields.**

The first flow shape is a doublet model, computed as $u = \frac{-2x_{\rm d}z_{\rm d}}{(x_{\rm d}^2+z_{\rm d}^2)^2}$ and $v = \frac{(-z_{\rm d}^2+x_{\rm d}^2)}{(x_{\rm d}^2+z_{\rm d}^2)^2}$ Because of the quadrotor's proximity to the ground in our experiments, I also applied a "doublet plus mirror" model, which adds the flow from the original doublet to the flow from another equal-magnitude doublet that has been mirrored around the ground plane, thus achieving zero flux in airflow at the ground [23]. To compute the mirrored flow, I fixed the hovering quadrotor at 2 m above the ground (to match the experimental altitude of the hovering quadrotor during trials when the moving quadrotor flew beneath it). The two doublet models have vertical and horizontal dependence and components.

I also simulated several flow fields with only horizontal flow dependence. The simplest of these was a uniform downward gust below the hovering quadrotor with width w equivalent to the rotor tip-to-tip distance of 0.54 m. Another model included a downward gust with a magnitude determined by a Gaussian curve centered at the hovering quadrotor's midpoint, $v = -\exp\frac{x_{\rm d}^2}{w^2}$. A variation on this model summed two Gaussians located at equal distances ($d/2$) away from the vehicle's center to simulate the downward jets found beneath each rotor.

The last flow model tested used Particle Image Velocimetry (PIV) data from prior work [9] to estimate the velocities below a hovering quadrotor. The flow field measurements were taken around a Crazyflie micro quadrotor, thus I scaled the position of all the measurements by the relative arm length of the hummingbird compared to the Crazyflie. The cross-section of data was taken at half an arm length below the vehicle. I considered only the vertical component of the flow, thus treating the PIV-informed model similarly to the vertically-independent functions described above.

The PIV flow was also normalized by its maximum value for equal comparison to the other flow models, i.e. the maximum value of the normalized flow is 1. The other models are normalized as follows: the maximum value 1 m below the doublet is 1, the maximum value of the Gaussian is 1, and the uniform flow has a magnitude of 1. The maximum values of the doublet plus mirror and double Gaussian models are not exactly 1 because of the interference of the second flow. Their maximum values are 0.94 and 1.02, respectively.

While the six flow models describe the general shape of the flow field, the magnitude of each flow must be scaled to better approximate real velocities beneath a hovering quadrotor. The baseline simulation used an anemometer measurement beneath the hovering hummingbird quadrotor to do this: I multiplied the normalized flow field velocities by our measured value of 6 m/s for a more realistic estimate of the flow's strength. The velocity was then used to determine $F_{\mathrm{gust},z}$ by approximating the drag on the quadrotor using $F_d = \frac{1}{2} A C_d \rho (6v)^2$, where C_d is drag coefficient and A is projected area. For the purposes of this heuristic ballpark estimate, I modeled the drag on the quadrotor as a half sphere (C_d = 0.42) with an area A calculated as a circle with a diameter $2d$. This forcing was then input into the force balance equations as $F_{\mathrm{gust},z}$.

The experimental findings informed my simulation process. First, we observed no deflection in the quadrotor's path when flying above another quadrotor, so I decided to only simulate paths below the hovering vehicle. The data also show no dependence in the real quadrotor's path deflection on its altitude beneath the hovering quadrotor (the variations between path deflections were negligible compared to the magnitude of those deflections). When comparing to our simulations, I therefore averaged the path data of all the altitudes at each speed. To address the altitude input required by our two doublet flow models, I chose to test a commanded relative altitude of 1 m across every test case. No modifications were needed for the other flow models since they have no vertical variation in the flow.

I included an additional constant upward flow velocity to all simulations. This proved necessary because the flight path of the actual quadrotor initially trended above its desired altitude throughout the entire parameter sweep, and this trend was not otherwise captured by our models. To determine the desired offset, I calculated the average boost in altitude of the experimental quadrotor during the middle half (x positions of -1.5 to -0.5 m) of it's approach toward the hovering quadrotor. I then calculated the upward flow velocity necessary for our simulated quadrotor to achieve the offset. This process was computed separately for slow (0.3 m/s) flights and fast (0.8 m/s) flights. The computed velocity was added to the flow at all locations throughout every simulation at each speed.

Because the values of the heuristically-determined parameters in the baseline model could impact the performance of the different flow models, I re-ran the simulations with several different levels of parameter fitting, using the root-mean-square error (RMSE) between the simulated and real altitudes at each sampled x position (-2 – 2 m in 0.01 m increments) as a measure of model fit. First, I replaced the heuristic drag coefficient with a fitted value. Next I iterated over both the strength and the width of several of our flow models, altering the value of w used to compute the Gaussian, double Gaussian, and uniform jet flow fields. Last, I used a gradient-descent algorithm to optimize the translational gains ($K_{\mathrm{P},x}$, $K_{\mathrm{D},x}$, $K_{\mathrm{P},z}$, $K_{\mathrm{D},z}$) in addition to the flow strengths and widths. This process identified a combination of parameters that minimized the RMSE locally to the initial heuristic values.

3.3 Experimental findings and modeling implications

The experimental data reveal several notable trends (Figure 3.5). When the moving quadrotor flew above the hovering quadrotor, its flight closely adhered to the commanded flight path. This is in contrast to the flights below the hovering quadrotor, where the vehicle was deflected downward by 7 – 8 cm. It appears that the flow above our quadrotor is characterized by negligibly small velocity gradients, meaning a different (or no) flow model is necessary for above-hovering flight.

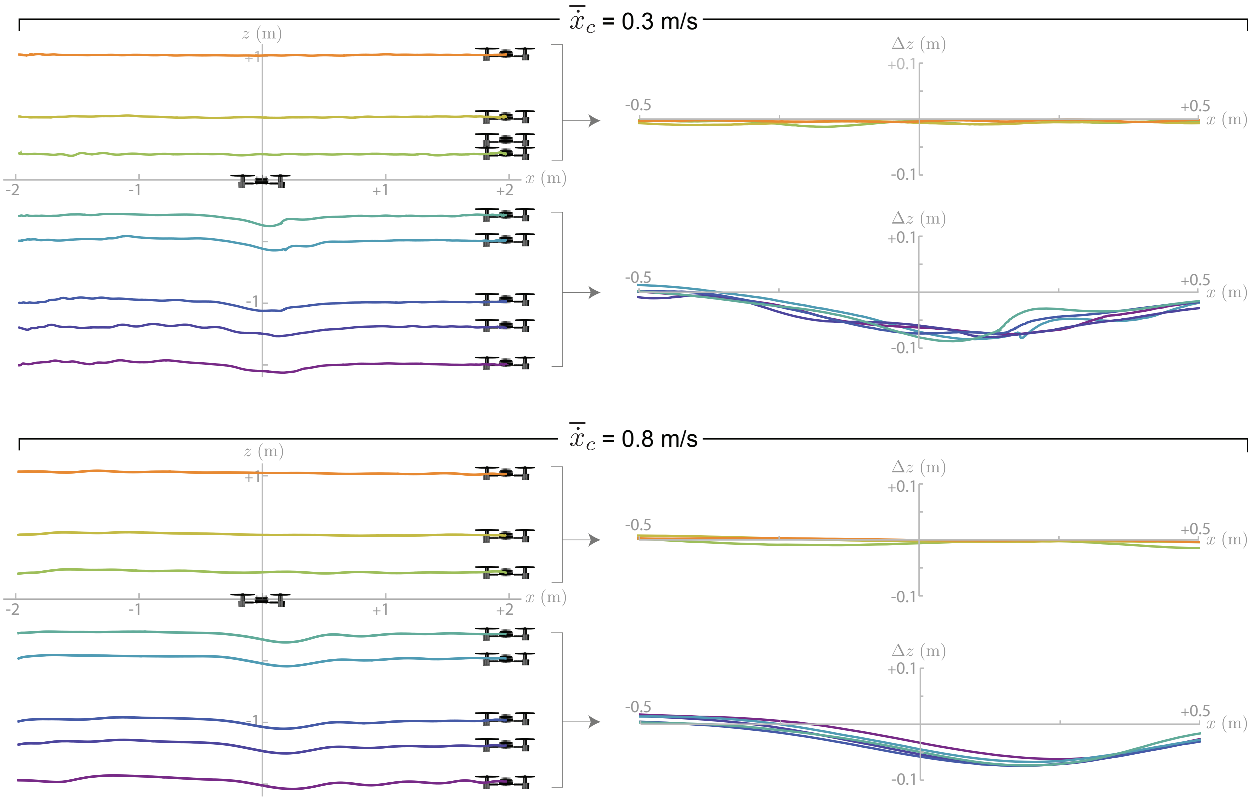

Figure 3.5: **Experimental Results. Left: flight paths relative to the hovering quadrotor. Right: flight paths relative to their commanded altitude.**

Of the below-hovering flights, there is no discernible pattern in the effect of relative commanded altitude on the magnitude of deflections; deflections did not vary significantly between flights as close as 0.3 m and as far as 1.5 m below the hovering quadrotor. It appears that the flow beneath is a strong jet that stays stable and coherent for several (7+) rotor diameters. Further experimentation is needed to understand the range of speeds and proximities that do significantly impact a quadrotor's flight path.

The absence of relative altitude dependencies (within the ranges tested) in the quadrotor's deflection has important modeling implications; effective models can be independent of these variables. A two-dimensional flow model, such as a doublet, may therefore be unnecessarily complex for simple disturbance-rejection path-planning. Our simulation results support this idea—our doublet models did not perform better than the dimensional models at predicting the quadrotor's deflection. Additionally, the absence of a disturbance when flying above a hovering vehicle suggests that no flow needs to be considered when flying in this region.

Our experimental data also show a systematic flight disturbance at the beginning of each flight: the moving quadrotor experiences an upward displacement of about 2-4 cm at distances far from the hovering aircraft. One possible explanation is that air circulation caused by the hovering quadrotor results in an updraft at distances far from the vehicle. While a doublet flow model does contain this type of updraft, our doublet simulations did not recreate the updraft seen in the experiments. Alternatively, the boost in altitude could be due to interactions with other surfaces in our indoor test environment; near- ground or ceiling flight is known to cause an increase in lift [9, 7, 5]. Simulations using our doublet plus mirror model, which approximates doublet flow near the ground, did not capture the experimental boost in altitude, however. Regardless of the unknown source of this disturbance, our conclusions about relative path deflection should not be affected by the upward displacement since we added an upward flow velocity to recreate this altitude offset.

3.4 Simulation Results

The simulated flights are depicted in Figure 3.6. The computed error for each simulation is shown in Figure 3.7. I first consider the results of the heuristic baseline simulations (Figure 3.6a). All the simulations were successful at predicting a downward deflection as the moving quadrotor passed beneath the hovering quadrotor. The simulation with the PIV-based flow field was the most successful at predicting the actual flight path while the other models predicted vertical deflections several times larger than the actual deflection. In the slow case, the simulated flights also overshot the target altitude after they recovered from the disturbance; it appears that the system is underdamped.

The results from the flow strength fitting are shown in Figure 3.6b. At each commanded speed, the simulated paths are similar across flow models. This uniformity suggests that the PIV-informed flowfield outperformed the other flow models in the baseline simulation not because of its shape but because its net forcing. Of the strength fit simulations, the slower trials more closely achieved a vertical deflection below the hovering quadrotor with magnitudes of about 4 cm. These underestimated the real deflection and, like in the heuristic simulation, overshot the target as they recovered from the disturbance.

Fitting the flow field width in addition to the strength (Figure 3.6c) resulted in little improvement in flight path compared to the simulations that only fit the flow strength. Because the strength and width were fitted together, we cannot conclude that the width or shape of the flow is inconsequential to the simulated flight path, but it appears that the net forcing of each shape is a determining factor to the accuracy of the simulated flight path.

The baseline model and first two fits predict a deflection that is delayed compared to the real deflection. This discrepancy improved when gains were optimized (Figure 3.6d). The deflection magnitude also aligned more closely with the experimental results; the simulations nearly achieved the experimental flight path. This improvement could be attributed to the reduction in oscillations after the quadrotor passes through the flowfield - the system no longer appears to be underdamped. This increased flight path accuracy is reflected in the reduced errors across all flow models shown in Figure 3.7.

The progression of model fits provides guidance for future modeling efforts. When the flow strengths are fitted to reduce the error in the flight paths, there is little difference in the trajectories between different flow models. This suggests that the average magnitude of the flow matters more to the result than does the shape of the flow. Additionally, the large increase in accuracy achieved by fitting the gains indicates that a vehicle's control is key to matching real flight behavior. Future work could seek to analytically explore the relationship between the control, flow field, and physical parameters, as this simulation can be collapsed into several dimensionless groups.

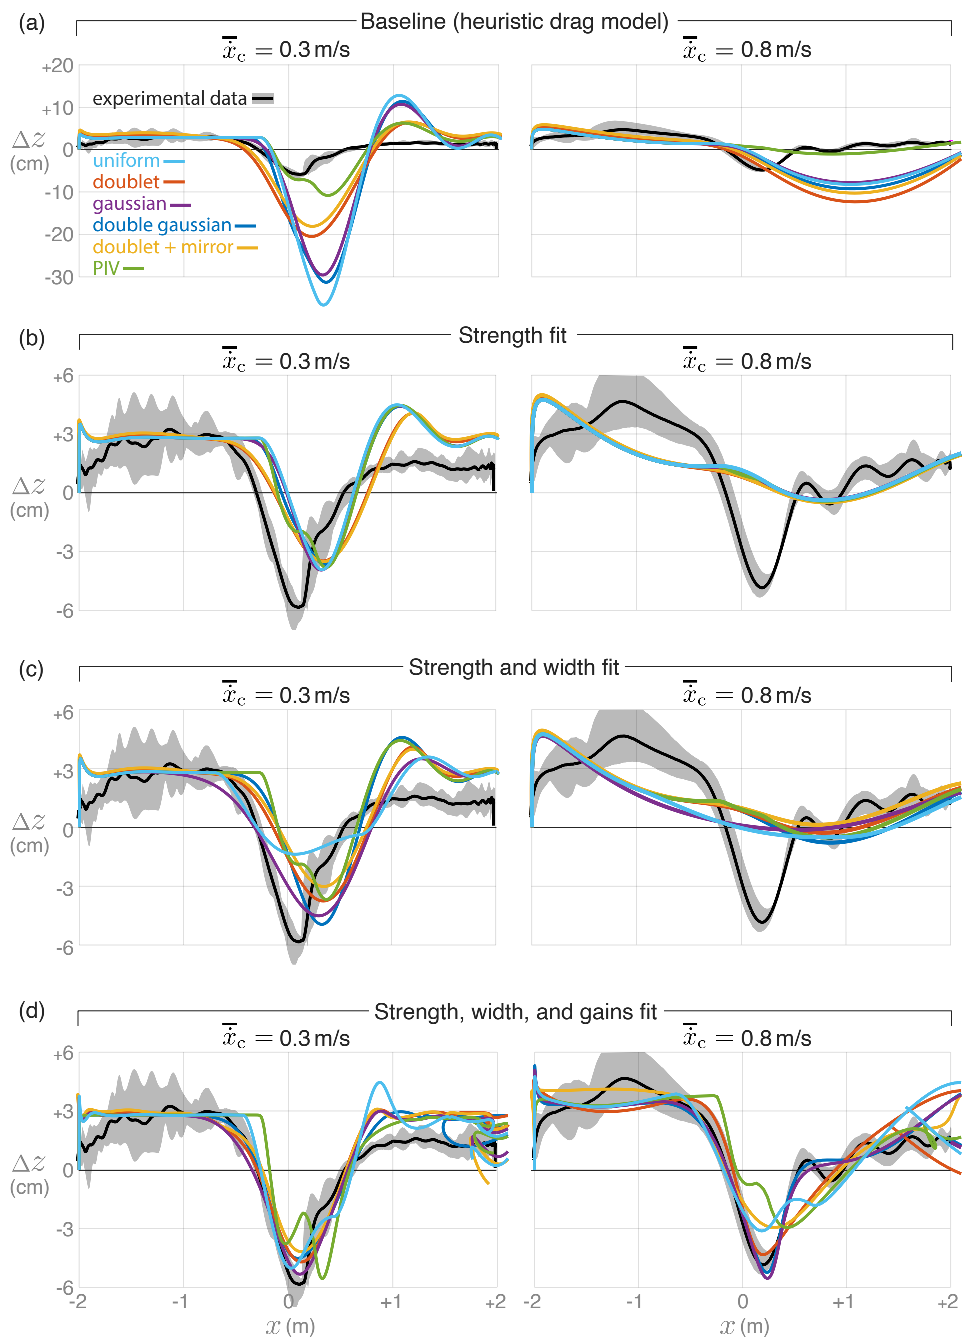

Figure 3.6: **Simulated flight paths at the faster of the two flight velocities. The averaged experimental data is shown for comparison. a-d show the progression of model fits from the heuristic baseline simulation (a) to the gains and flow-fitted simulation (d).**

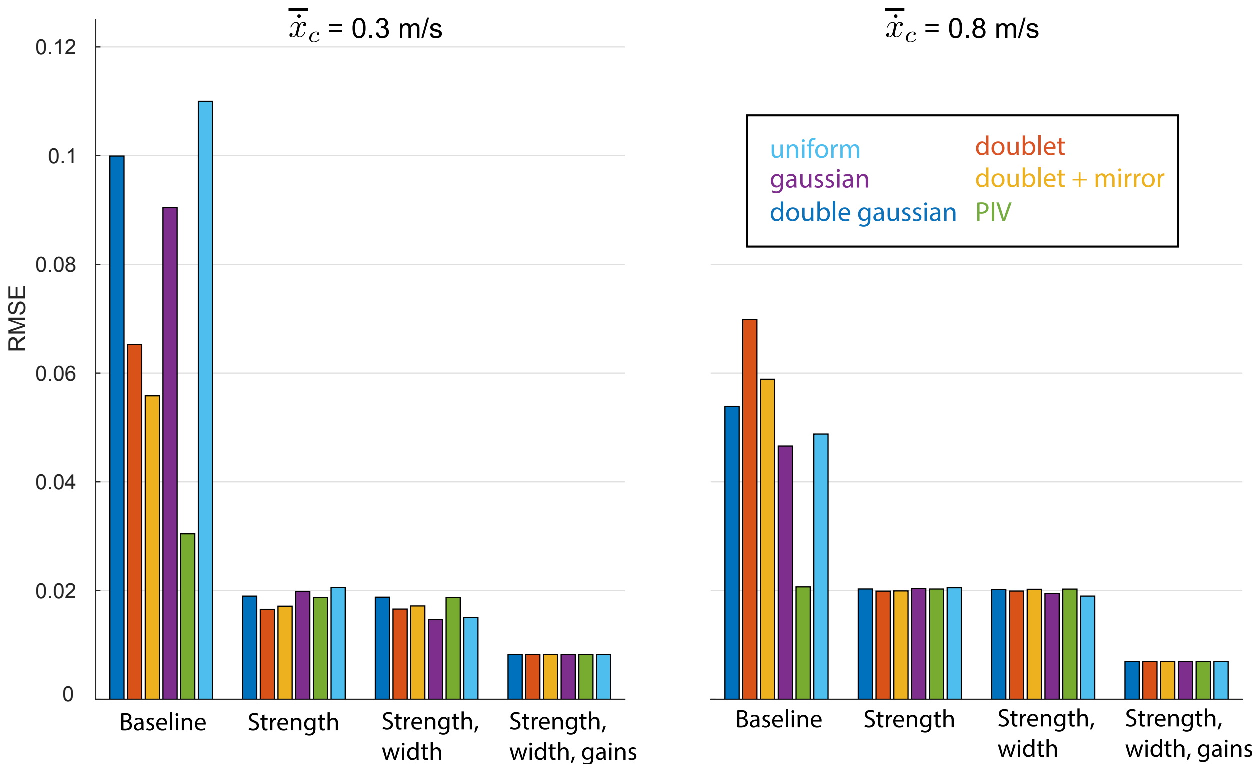

Figure 3.7: **The Root-Square Mean Error of the simulated quadrotor's flight path in predicting the experimental flight path. Each bar represents one combination of flight velocity, flow model, and level of fit; all simulations are summarized here.**

4 Conclusion

4.1 Key Findings

The 1D near-boundary simulations (Chapter 2) show the importance of accounting for the boundary effect when planning a flight path that both optimizes power consumption and does not present a significant risk of crashing. Simulations indicate that a quadrotor can safely fly closer to the ground than would be expected had the increased lift not been accounted for (Figure 2.1a,c). In contrast, flights must stay further from the ceiling to ensure the same level of safety that would be expected without considering the increased lift near the ceiling (Figure 2.1b,d).

The near-boundary simulations also give insight into the importance of the relative magnitude of the boundary effect. The safe flight altitude is more sensitive to the magnitude of the ceiling effect than it is to the magnitude of the ground effect (Figure 2.1c-d). This indicates that researchers need a higher degree of confidence in their ceiling model in order to safely fly within this region ($1.5 \lesssim \hat{z} \lesssim 0.5$). On the other hand, because the flight can occur closer to the ground ($0.75 \lesssim \hat{z} \lesssim 1.25$) where the energy cost curve is the steepest, even small decreases in ground altitude can result in large efficiency improvements. For that reason, researchers might be interested in further refining their understanding of the ground effect, especially at distances below $\hat{z} \approx 0.5$, where experimental results begin to deviate from current models [9].

The quadrotor-quadrotor disturbance simulations (Chapter 3) were partially successful at predicting the deflections in the moving quadrotor's flight associated with the downwash of the stationary quadrotor. The heuristic baseline model, which was designed to use only easy-to-estimate parameters and required no model fitting, over-predicted the magnitude of deflection in all but one case (using the PIV data to predict a 0.8 m/s flight). Because most use cases are concerned with ensuring the quadrotor does not deviate too far from its commanded trajectory and, for example, crash into an obstacle close below it, an over prediction could be acceptable as it allows a factor of safety in path-planning. Experiments with a variety of quadrotor sizes and flight conditions are needed to confirm the efficacy of this method in estimating the magnitude of deflection, but this approach appears to work well for the specific quadrotors we used. Fitting to the average magnitude of the flow and/or the quadrotor gains will result in an improvement in model predictions.

Finally, this study as a whole presents a novel method to learn about complex fluid-structure-control interactions using relatively simple flow simulations and experiments. In situations without a measurement of the flow field and where a detailed calculation of the flow would be computationally expensive, a researcher could apply a simple flow model (or multiple superimposed flow models, as in the panel method [24]) or boundary model within a dynamic simulation to estimate the flow interactions. Iterative comparisons to experimental data could reveal valuable insights about the flow field itself as well as the system as a whole. I've shown this strategy to be applicable to quadrotor flight near boundaries and near another quadrotor.

4.2 Contributions

This work adds flow models to rigid-body dynamical simulations of quadrotor flight. The simulation results contain valuable insight into the importance of local flow on a quadrotor during two specific near-obstacle flight scenarios. In producing these results, this work illustrates the use of this novel modeling technique to UAV system operators and researchers.

Chapter 2, published as the numeric portion of an article in the American Institute of Aeronautics and Astronautics (AIAA) Journal [9], outlines a 1D numeric simulation that accounts for the increased lift as a result in the ground or ceiling effect. The simulation illustrates the tradeoffs between safety and efficiency and can provide a guide for quadrotor users to choose a balance that suits their particular UAV application. The results also indicate the importance of a correct estimation of the boundary effect, especially when flight approaches a ceiling.

Chapter 3, which is under review in the AIAA Journal [25], explores close range flight between two quadrotors via a novel simulation technique that combines rigid-body dynamics and canonical flow models. I show that this technique can be useful as a simple heuristic method to roughly estimate flight path disturbance beneath a hovering quadrotor. The study also shows how the process of fitting and refining this type of reduced-order model can itself give researchers valuable insight into the form of the real flow. This could streamline an experimental process, which is otherwise cumbersome as a result of the large number of factors influencing quadrotor-quadrotor flight dynamics.

4.3 Future Work

My simulations in Chapter 2 illustrate the importance of an accurate understanding of a boundary model on safely taking advantage of the energy savings associated with near-boundary flight. Given this, a logical next step is to further refine our understanding of near-boundary flight. Work is currently underway at UVA MAE to characterize the impact of a rotor's blade twist and size on near-surface control dynamics. I designed significant portions of the initial test rig (Figure 4.1) in collaboration with Darius Carter and Qiang Zhong, who have since produced an improved version. The trends revealed from this study can be incorporated into a similar reduced-order simulation similar to the one used in Chapter 2 to explore the relationship between rotor shape and crash propensity.

Other work might directly build off of my simulations from Chapter 3. It would be interesting to add multiple potential flows (as in the panel method) and/or an analytic ground effect model to the

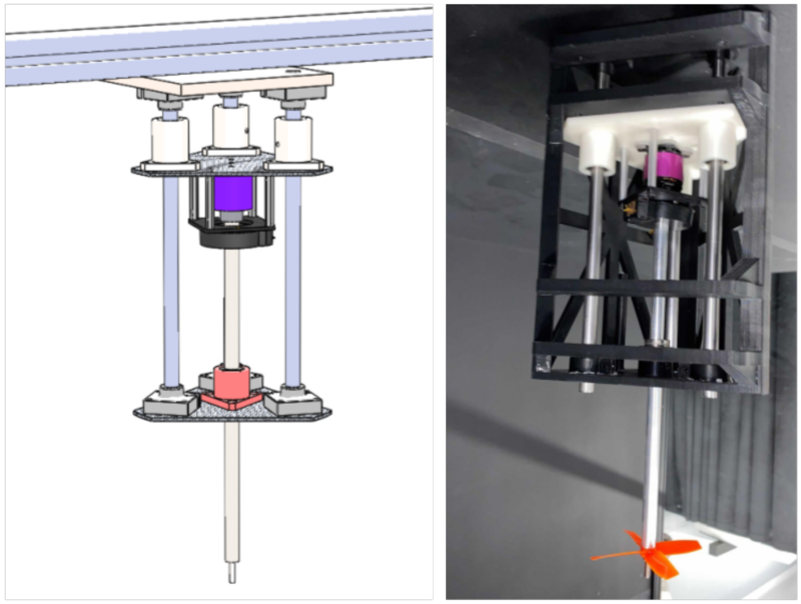

Figure 4.1: **Prototypes of a test setup designed in collaboration with Darius Carter. The central axle, the motor (in purple), and the propeller (in orange in the right image) can move vertically in response to changes in the local fluids or in motor power. Photograph provided by Darius Carter.**

reduced-order simulation presented here in an attempt to improve the accuracy of the simulated flight path. UAV operators might also want a wider characterization of quadrotor-quadrotor flow interactions and seek to expand upon the parameters tested in this study. Ultimately, UAV operators could use these reduced-order simulations to design a controller that compensates for the path deflection due to a nearby quadrotor's downwash.

Bibliography

[1] B. Terwilliger, D. Ison, J. Robbins, and D. Vincenzi, *Small Unmanned Aircraft Systems Guide: Exploring Designs, Operations, Regulations, and Economics*. Aviation Supplies & Academics, Incorporated, 2017, ISBN: 9781619543942.

[2] B. Barrett, *Inside the olympics opening ceremony world-record drone show*, 2018.

[3] H. Nobahari and A. Sharifi, "Continuous ant colony filter applied to online estimation and compensation of ground effect in automatic landing of quadrotor," *Engineering Applications of Artificial Intelligence*, vol. 32, pp. 100–111, 2014.

[4] L. Danjun, Z. Yan, S. Zongying, and L. Geng, "Autonomous landing of quadrotor based on ground effect modelling," in *2015 34th Chinese Control Conference (CCC)*, IEEE, 2015, pp. 5647–5652.

[5] S. Gao, C. Di Franco, D. Carter, D. Quinn, and N. Bezzo, "Exploiting ground and ceiling effects on autonomous uav motion planning," in *2019 International Conference on Unmanned Aircraft Systems (ICUAS)*, IEEE, 2019, pp. 768–777.

[6] P. Sanchez-Cuevas, G Heredia, and A Ollero, "Multirotor uas for bridge inspection by contact using the ceiling effect," in *2017 International Conference on Unmanned Aircraft Systems (ICUAS)*, IEEE, 2017, pp. 767–774.

[7] P. Sanchez-Cuevas, G. Heredia, and A. Ollero, "Characterization of the aerodynamic ground effect and its influence in multirotor control," *International Journal of Aerospace Engineering*, vol. 2017, 2017.

[8] Y. H. Hsiao and P. Chirarattananon, "Ceiling effects for hybrid aerial-surface locomotion of small rotorcraft," *arXiv preprint arXiv:1905.04632*, 2019.

[9] D. J. Carter, L. Bouchard, and D. B. Quinn, "Influence of the ground, ceiling, and sidewall on micro-quadrotors," *AIAA Journal*, pp. 1–8, 2020.

[10] S. L. Waslander, G. M. Hoffmann, J. S. Jang, and C. J. Tomlin, "Multi-agent quadrotor testbed control design: Integral sliding mode vs. reinforcement learning," in *2005 IEEE/RSJ International Conference on Intelligent Robots and Systems*, IEEE, 2005, pp. 3712–3717.

[11] B. T. Clough, "Uav swarming? so what are those swarms, what are the implications, and how do we handle them?" AIR FORCE RESEARCH LAB WRIGHT-PATTERSON AFB OH AIR VEHICLES DIRECTORATE, Tech. Rep., 2002.

[12] P. Vincent and I. Rubin, "A framework and analysis for cooperative search using uav swarms," in *Proceedings of the 2004 ACM symposium on Applied computing*, 2004, pp. 79–86.

[13] Y. Altshuler, V. Yanovsky, I. A. Wagner, and A. M. Bruckstein, "Efficient cooperative search of smart targets using uav swarms," *Robotica*, vol. 26, no. 4, p. 551, 2008.

[14] V. Sharma, D. N. K. Jayakody, I. You, R. Kumar, and J. Li, "Secure and efficient context-aware localization of drones in urban scenarios," *IEEE Communications Magazine*, vol. 56, no. 4, pp. 120–128, 2018.

[15] E. Barmpounakis and N. Geroliminis, "On the new era of urban traffic monitoring with massive drone data: The pneuma large-scale field experiment," *Transportation research part C: emerging technologies*, vol. 111, pp. 50–71, 2020.

[16] I. Cheeseman and W. Bennett, "The effect of the ground on a helicopter rotor," *R & M*, vol. 3021, 1957.

[17] S. Bouabdallah and R. Siegwart, "Full control of a quadrotor," in *2007 IEEE/RSJ International Conference on Intelligent Robots and Systems*, Ieee, 2007, pp. 153–158.

[18] Y. Song, S. Naji, E. Kaufmann, A. Loquercio, and D. Scaramuzza, "Flightmare: A flexible quadrotor simulator," *arXiv preprint arXiv:2009.00563*, 2020.

[19] D. Mellinger and V. Kumar, "Minimum snap trajectory generation and control for quadrotors," in *IEEE International Conference on Robotics and Automation*, 2011, pp. 2520–2525.

[20] E. Yel, T. X. Lin, and N. Bezzo, "Computation-aware adaptive planning and scheduling for safe unmanned airborne operations," *Journal of Intelligent & Robotic Systems*, vol. 100, no. 2, pp. 575–596, 2020.

[21] C. S. Subudhi and D Ezhilarasi, "Modeling and trajectory tracking with cascaded pd controller for quadrotor," *Procedia computer science*, vol. 133, pp. 952–959, 2018.

[22] J. G. Ziegler, N. B. Nichols, *et al.*, "Optimum settings for automatic controllers," *trans. ASME*, vol. 64, no. 11, 1942.

[23] J. Katz and A. Plotkin, *Low-speed aerodynamics*. Cambridge university press, 2001, vol. 13.

[24] L. L. Erickson, "Panel methods: An introduction," 1990.

[25] L. Bouchard, E. Yel, N. Bezzo, and D. B. Quinn, "Quadrotor-quadrotor interactions: Experiments and potential-flow-informed simulations of flights through another vehicle's wake," *AIAA Journal*, Under Review.