Anemoi: A Low-cost Sensorless Indoor Drone System for Autonomous Mapping of 3D Airflow Fields

Stephen Xia[1†*], Minghui Zhao[1†], Charuvahan Adhivarahan[2*], Kaiyuan Hou[1],
Yuyang Chen[2], Jingping Nie[1], Eugene Wu[1], Karthik Dantu[2], Xiaofan Jiang[1]
[1]Columbia University, [2]University at Buffalo
{stephen.xia, mz2866, kh3119, jn2551}@columbia.edu, ewu@cs.columbia.edu, jiang@ee.columbia.edu
{charuvah, yuyangch, kdantu}@buffalo.edu

ABSTRACT

Mapping 3D airflow fields is important for many HVAC, industrial, medical, and home applications. However, current approaches are expensive and time-consuming. We present Anemoi, a sub-$100 drone-based system for autonomously mapping 3D airflow fields in indoor environments. Anemoi leverages the effects of airflow on motor control signals to estimate the magnitude and direction of wind at any given point in space. We introduce an exploration algorithm for selecting optimal waypoints that minimize overall airflow estimation uncertainty. We demonstrate through microbenchmarks and real deployments that Anemoi is able to estimate wind speed and direction with errors up to 0.41 m/s and 25.1° lower than the existing state of the art and map 3D airflow fields with an average RMS error of 0.73 m/s.

CCS CONCEPTS

• **Computer systems organization → Sensors and actuators**; **Robotic autonomy**; *Sensor networks.*

KEYWORDS

sensor, sensorless, drone, micro aerial vehicles, mobile sensing, environmental sensing, path planning

ACM Reference Format:

Stephen Xia, Minghui Zhao, Charuvahan Adhivarahan, Kaiyuan Hou, Yuyang Chen, Jingping Nie, Eugene Wu, Karthik Dantu, Xiaofan Jiang. 2023. Anemoi: A Low-cost Sensorless Indoor Drone System for Autonomous Mapping of 3D Airflow Fields. In *The 29th Annual International Conference on Mobile Computing and Networking (ACM MobiCom '23), October 2–6, 2023, Madrid, Spain.* ACM, New York, NY, USA, 16 pages. https://doi.org/10.1145/3570361.3613292

1 INTRODUCTION

Maintaining good airflow in indoor spaces such as homes and offices is important to keep them well ventilated, improve overall air quality and comfort of the occupants. A well-ventilated room frequently removes stale air, humidity, and indoor pollutants (dust, particles, viruses/bacteria, etc.) which also benefits the respiratory health of the occupants. This is especially important in preventing the spread of diseases such as COVID-19 [34, 35, 39] and other airborne or moisture-based viruses and bacteria that could linger indoors for several hours if proper ventilation is lacking. Correspondingly, accurate measurement of airflow over time in a building would allow homeowners and building managers to optimize ventilation to eliminate “dead-zones” where little air is moving or replaced, while also limiting HVAC use in zones that do not need them [17]. Another application that would benefit from this knowledge is deciding on suitable locations to place air purifiers to circulate air, as well as remove dust and other particles. A third motivating example is that a building’s HVAC system is spec’d and provisioned before the building is populated with furniture and occupants. The dynamics of air that moves through the building will change significantly over time as new occupants, furniture, and other materials move into the space. Therefore, periodic airflow measurements are necessary for proper reconfiguration of the overall ventilation.

However, it can be difficult and expensive to measure the full airflow over an entire space. Typically, a professional walks through the building with an airflow measurement device to measure *airflow vectors* (direction and speed) at different locations, which is costly and labor-intensive [42].

---

†These authors contributed equally to this work.
* Corresponding authors.

ACM MobiCom '23, October 2–6, 2023, Madrid, Spain

ACM ISBN 978-1-4503-9990-6/23/10…$15.00
https://doi.org/10.1145/3570361.3613292

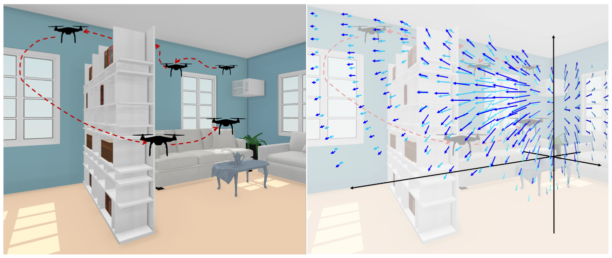

Figure 1: Illustration of mapping 3D airflow fields.

These measurements are aggregated over each location to obtain the full *airflow field* of the entire space, as shown in Figure 1. Another solution is to deploy static flow sensors throughout the environment. However, this can only generate accurate estimates for small enclosed spaces, such as within a fume hood. Our vision is to provide an inexpensive and automated process for measuring airflow gradients in 3D within large indoor environments using inexpensive drones. We believe this idea is timely with the introduction of small drones in indoor spaces like Ring Always Home Cam [56].

One solution is to attach a flow sensor to a drone and take measurements. However, this approach adds cost and significant weight, reducing the battery life and stability of the drone. For the safety, it is not feasible to fly large drones that may be more able to carry the weight of a flow sensor. We propose Anemoi, a drone-based system for autonomously mapping 3D airflow fields indoors without any additional sensors and only uses the sensors and control signals available on a typical micro-air vehicle. The key idea behind Anemoi is that the force from the wind flow will slightly perturb the drone, causing the onboard flight controller to compensate for the observed changes in position and orientation by actuating its motors. Anemoi leverages how the drone controller compensates for the force caused by the external wind to estimate airflow vectors (wind speed and direction) at a specific location in 3D space. We compare against existing state-of-art drone-based methods for measuring airflow that use physical models of the drone and directly leverage onboard sensors, like optical flow, and show that Anemoi achieves 0.4m/s and 25° lower wind speed and direction estimation error. It follows that the full airflow field can be measured by actuating the drone to each location and measuring the airflow vector. However, measuring every point in space can be time-consuming and inefficient, as small indoor drones can often fly for a few minutes before needing a recharge. Instead, we propose an information theoretic method to identify locations to sample airflow vectors that minimize the overall uncertainty of the current airflow field estimate, while minimizing the total travel distance of the drone. Next, we leverage *computational fluid dynamics* (CFD) to extrapolate and estimate the entire continuous airflow field from the airflows sampled at discrete locations. Through deployments in a real setting, we show that Anemoi can achieve higher accuracy estimates of full 3D airflow fields with fewer locations sampled compared to other sampling strategies.

While our system is evaluated to be suitable for airflow analysis for ventilation, we believe that there are a multitude of important applications of micro-drones that are orthogonal to our application (e.g., security and monitoring [57], asset inspection [44], and search and rescue [23]) that will spur an increase in intelligent micro-drones in the future. We envision that Anemoi will be one of the many important applications and services that future intelligent indoor micro-drones will provide. This work is an extended version of our demo [37]. Our contributions are as follows:

- We create Anemoi, a drone-based system for mapping 3D airflow vectors without using additional sensors by leveraging perturbations observed by the drone through its onboard feedback and flight controller to accurately measure 3D airflow vectors at a given location. We show that Anemoi achieves angular error of 6.20° and speed error of 0.11m/s, in comparison to 31.30° and 0.52m/s to a baseline method that requires the use of onboard sensors.
- We propose an information-theoretic exploration algorithm for the UAV to autonomously select 3-D points to minimize the overall estimation uncertainty of the full airflow field while also minimizing total distance traveled. Next, we leverage *computational fluid dynamics* to extrapolate the full 3D airflow field from 3D airflow vectors measured at individual sample points.
- We demonstrate through detailed simulations and real deployments that Anemoi requires traveling only half the distance to collect as many samples compared to a random walk and estimates airflow in a target space with 0.73m/s RMS error and average angle error of 29.79°.

2 RELATED WORK

To traditionally sense airflow vectors and fields, a person or professional would either need to move around the space with airflow sensors [21, 29], or install dedicated and stationary hardware (e.g., particle image velocimetry [14]) and using computational fluid dynamics [78], which are expensive and time-consuming. There are several existing works that leverage a drone or robot to measure airflow, including solutions that attach a wind sensor directly onto the drone/robot [40, 59, 69]. Other works use sensors on the drone or attach other sensors, such as an optical flow, intertial measurement unit (IMU), or camera sensor [46, 50, 58, 62, 75]. Instead, our work does not directly use any particular sensor, but rather the control signals sent by the drone's controller to actuate its motors to react to perturbations caused by airflow. This means that Anemoi can be adapted to any standard drone that performs self-stabilization, regardless of the sensors it has onboard, unlike existing works.

Other methods for monitoring airflow and the condition of the atmosphere (e.g., turbulence) leverage doppler radars [10, 18, 72], lidars [24, 32, 60, 77] and sodars [6, 26, 31, 36, 49]. All methods estimate wind speed similarly by measuring the doppler effect of transmitted and reflected signals caused by airflow and turbulence. The main difference between each sensor is the type of signal that is transmitted (radar = radio waves, lidar = light/laser, sodar = acoustic waves). All of these systems are used to measure outdoor atmospheric conditions, where wind speeds are orders of magnitude greater (tens to hundreds of m/s compared to < 5 m/s indoors). Additionally, the observed error is often on the order of several m/s, which is too high for indoor airflows that may only reach several m/s. Moreover, these sensors are too large for a small drone to carry, often weighing more than tens of kilograms [51].

Anemoi measures changes in motor control signals transmitted by the flight controller, which is used to correct for any position, velocity, and attitude errors. Most personal and open-source drones integrate model free proportional-integral-derivative (PID) controllers, which employ feedback from sensors to continuously adjust motors to maintain position, velocity, and stability [1, 3]. There are a number of other control schemes that model flight dynamics, such as non-linear dynamic inversion (NDI) and incremental non-linear dynamic inversion (INDI) [64–66, 80], reactively adjust motors when errors or instability naturally becomes too great [9], or incorporate mission objectives into the control loop through reinforcement learning [48, 74]. For this paper, we only look at the performance of Anemoi under traditional PID control and leave other schemes to future work.

Exploration [11, 63, 71] is a well studied field in robotics. In [71] the authors use visual sensors to build and explore on a 2D map of the environment. This idea of exploration has been applied to other problems such as 3D reconstruction [30], active tactile sensing [47], etc. In this work, we adapt information theoretic robotic exploration to efficiently sense airflow. Path planning algorithms [55] find navigable paths avoiding obstacles. Active mapping or exploration algorithms [52] find efficient navigable paths (planning) with the joint objective of improving maps (mapping) built through sensing the environment during navigation. While other methods perform active mapping to improve visual maps [52], we employ these techniques to perform better airflow estimation in the target space.

3 ESTIMATING 3D AIRFLOW GRADIENTS

We intend to use small drones (i.e., micro-air vehicles) to sense airflow, since these will be commonplace in indoor spaces with the advent of devices such as the Ring Always Home Cam [56]. Our objective is to utilize the effect of airflow on control signal variations of an autonomously hovering drone to characterize the direction and strength of airflow at any given point. In this section, we propose a method to estimate airflow at a single point. In section 4 we introduce a method to efficiently sample and estimate the airflow field in an entire target space.

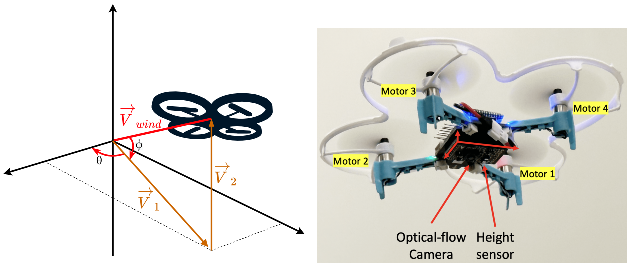

Figure 2: Quantities to estimate airflow vectors. (Left: $\vec{V}_{wind}$, $\vec{V}_1$, and $\vec{V}_2$ are the 3D wind vector, the azimuthal-plane, and z-axis wind speed components. θ and ϕ are the azimuth and polar angle of the wind. Right: MiniFly drone with axes and motors labeled).

3.1 Effects of Wind on Motor Control Signals

To leverage motor control signals for airflow estimation, we first analyze the effects of airflow on these signals as the drone hovers at a particular *setpoint* location. As shown in Figure 2, we need to find $\vec{V}_{wind} : (|\vec{V}_{wind}|, \theta, \phi)$, which correspond to the wind speed, azimuth, and polar angle, respectively. Figure 3 also plots the motor control signals from the flight controller as the drone hovers in place. These values are directly proportional to the rotation speed of each motor. To make a quadrotor hover, the opposite motors spin at the same speeds, and the adjacent ones spin at the same speed but in opposite directions. Therefore, when there is no significant airflow, the signals sent to each motor are very similar in magnitude as seen in Figure 3.

When we blow a constant stream of air, we see that the control signals to the motors change significantly. The force from the airflow causes instability in the drone's roll and pitch that the drone can observe using its onboard sensors (IMU, height sensor, camera). Almost all drone systems utilize a proportional-integral-derivative (PID) based feedback controller to determine the correct control signals to send to the motors to maintain a stable pose in real time. Blowing air at $\theta = 0$ degrees and $\phi = 0$ degrees causes the motors 1 and 2 to spin faster and the difference between these two motors and motors 3 and 4 become significantly greater as the controller attempts to account for instability caused by the wind. When we increase the wind speed, we see that the difference in motor speeds becomes even greater. This suggests that **the magnitude in the differences between motor control signals inform wind speed**, $|\vec{V}_1|$. When we change the airflow direction to $\theta = 90$ degrees, we see

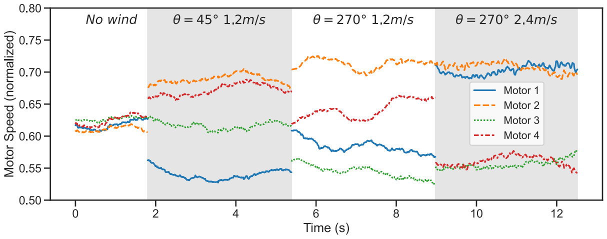

Figure 3: Control signals sent to each motor at different wind speeds and directions.

that the signals to motors 1 and 3, which are closer to $\theta = 90$ degrees than motors 2 and 4, increase significantly. This suggests that **the polarity of the differences in motor control signals can inform the direction of the wind**. Every drone may utilize different sensors to maintain stability, but it sends control signals to its motors to move and maintain stability. Utilizing motor control signals allows us to estimate airflow vectors without requiring a specific set of sensors, as the PID controller "abstracts" the sensors it uses, unlike existing methods. This is unlike existing methods that require specific types of sensors installed onto the drone (e.g., GPS [46] or optical flow [58]).

3.2 Estimating Airflow Vectors from Motor Control Signals

Estimating Azimuth Air Vector: Figure 4a shows the motor speed differentials between motors 1 and 2 ($MD_{12} = motor_1 - motor_2$), motors 1 and 3, and motors 1 and 4. We varied the direction of the airflow in eight equally spaced azimuth directions (every 45 degrees) around the drone. We also collect data for two different wind speeds in each direction. We describe details on how we designed and implemented a stable and repeatable wind generation for our model training and evaluation in Section 5.1. We see that the motor differentials tend to follow a single line if the wind is blowing from the same direction (Figure 4a). If we view the data from the plane normal to these lines (Figure 4b), we see that all points map onto a circular structure, which is physically intuitive. Additionally, we see that stronger winds tend to extend outward, creating a larger circle. The basis of our airflow estimation algorithm involves finding the projection that maps 3D motor differentials onto a 2D plane, such that the angle of observed motor differentials on this plane corresponds to the azimuth direction of the airflow, while the distance of an observation corresponds to the wind speed (Figure 4c and 4d).

Our goal is to project every observation corresponding to the same wind direction and speed to one point that estimates airflow. To start, we are given $\mathbf{X} = [\ \mathbf{x}_1\ \mathbf{x}_2\ \mathbf{x}_3\ \dots\ \mathbf{x}_n\ ]$ where $\mathbf{X}$ is a collected dataset of n motor differential vectors that we will use to learn the projection that maps every observation onto the direction of the wind. $\mathbf{x}_i = [MD^i_{12}, MD^i_{13}, MD^i_{14}]^T$ is the i-th three-element motor differential vector in the dataset, and $\mathbf{Y} = [\ \mathbf{y}_1\ \mathbf{y}_2\ \mathbf{y}_3\ \dots\ \mathbf{y}_n\ ]$ where $\mathbf{y}_i$ is the labeled direction and speed $(\theta_i, |\vec{V}^i_1|)$. With this dataset, we aim to find a projection $\mathbf{Q} : \mathbb{R}^3 \mapsto \mathbb{R}^2$ as shown in Equation 1.

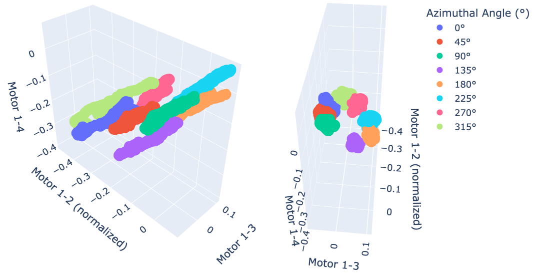

(a) Motor differential **(b) Cross-section of (a)**

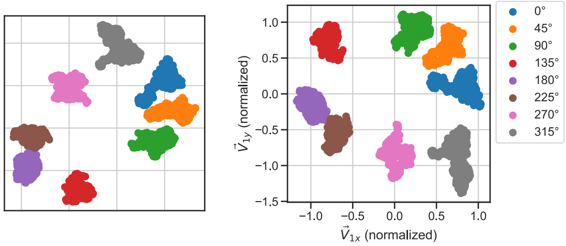

(c) Motor differentials using Q. **(d) Projection after transformed into a plane**

Figure 4: 3D visualization ((a) and (b)) and projections ((c) and (d)) of motor differentials.

$\overline{\mathbf{x}}_i$ is the centroid of class i, or the mean of all observations belonging to class i, where each class i is described by an azimuth and wind speed pair: $(\theta_i, |\vec{V}^i_1|)$. In other words, during the training or calibration process, wind is blown at the drone at a specific angle and wind speed and multiple samples of the observed motor differentials may be collected, and this process is repeated for d different $(\theta_i, |\vec{V}^i_1|)$ pairs. As an example, in Figures 4c and 4d, we collected $d = 8$ different classes, varying the azimuth direction of the wind equally around the drone. $\mathbf{c}^{proj}_i$ is the centroid of class i after being projected by $\mathbf{Q}$, while $\mathbb{1}(i, j)$ is an indicator function that equals 1 if $\mathbf{x}_i$ belongs to class j, $(\theta_j, |\vec{V}^j_1|)$, and 0 otherwise. The constraint $\mathbf{Q}\mathbf{Q}^T = I$ is in place to ensure that our solution yields a basis for a new coordinate system, where I is the identity matrix.

$$\begin{aligned}
&\arg\min_{\mathbf{Q}}(\mathbf{L}_P + \mathbf{L}_C) \\
&\mathbf{Q}\mathbf{Q}^T = I \\
&\quad \mathbf{L}_P = a\sum_{j=1}^{d}\sum_{i=1}^{n}\tfrac{1}{2}\|\mathbf{Q}\mathbf{x}_i - \mathbf{c}^{proj}_j\|_2^2\mathbb{1}(i,j) \\
&\quad \mathbf{L}_C = -b\sum_{j=1}^{d}\sum_{k=j+1}^{d}\|\mathbf{c}^{proj}_j - \mathbf{c}^{proj}_k\|_2^2 \\
&\mathbf{c}^{proj}_i = \mathbf{Q}\overline{\mathbf{x}}_i,\ i = 1, \dots, d
\end{aligned} \tag{1}$$

Additionally, the loss function that guides the search for our mapping function, $\mathbf{Q}$, is composed of two components:

$\mathbf{L}_P$ and $\mathbf{L}_C$. $\mathbf{L}_P$ finds the mapping such that all observations belonging to the same class j $(\theta_j, |\vec{V}_1^j|)$ are mapped to the same point by minimizing the distance of the projected observations to the centroid of class j. To increase sensitivity and regress between different wind angles and speeds, we want the classes on the projected plane to be as far away and distinct as possible. This involves increasing the distance between the centroids of each class, which is reflected in the $\mathbf{L}_C$. a and b are constants to weight the importance of each loss. In our experiments, we found that $a = 1.0$ and $b = 0.1$ yielded the best performing models.

To solve this optimization problem, we apply the method of Lagrange multipliers, as shown in Equation 2. Setting the gradient with respect to $\mathbf{Q}$ equal to 0 yields $\frac{\partial \mathbf{L}}{\partial \mathbf{Q}} = \mathbf{VQ} - \lambda \mathbf{Q} = 0$ which implies $\mathbf{VQ} = \lambda \mathbf{Q}$. This result implies that the optimal coefficients of λ are eigenvalues of $\mathbf{V}$ and that the bases that form our projection matrix $\mathbf{Q}$ are the corresponding eigenvectors. To find the final projection $\mathbf{Q}$, we compute and take the eigenvectors of the top two eigenvalues of $\mathbf{V}$.

$$\begin{aligned}
\mathbf{L} &= \mathbf{L}_P + \mathbf{L}_C - \lambda(\mathbf{Q}\mathbf{Q}^T - I) \\
\frac{\partial \mathbf{L}}{\partial \mathbf{Q}} &= \mathbf{VQ} - \lambda \mathbf{Q} \\
\mathbf{V} &= \textstyle\sum_{j=1}^{d} \sum_{i=1}^{n} \mathbf{x}_{ij}\mathbf{x}_{ij}^T \mathbb{1}(i,j) - b \sum_{j=1}^{d} \sum_{k=j+1}^{d} \mathbf{c}_{jk}\mathbf{c}_{jk}^T \\
\mathbf{x}_{ij} &= \mathbf{x}_i - \overline{\mathbf{x}}_j,\ \mathbf{c}_{jk} = \overline{\mathbf{x}}_j - \overline{\mathbf{x}}_k
\end{aligned} \tag{2}$$

Figure 4c shows the plot of motor differentials after applying projection $\mathbf{Q}$. We see that our observations have all been mapped closely, forming a circular shape. However, this circular pattern is slightly distorted, rotated, and translated from a circle that directly maps to the azimuth of the wind, θ. To correct for this, we learn a transformation $\mathbf{T} : \mathbb{R}^3 \mapsto \mathbb{R}^3$ shown in Equation 3.

$$\arg\min_{\mathbf{T}} \tfrac{1}{2} \textstyle\sum_{j=1}^{d} \|\mathbf{T}\mathbf{c}_j^{proj,h} - \mathbf{r}_j^h\|_2^2 \tag{3}$$

$\mathbf{c}_j^{proj,h}$ represents the centroid for class j $(\theta_j, |\vec{V}_1^j|)$ in homogeneous coordinates. Representing our observations in the homogeneous coordinate system allows us to learn translation transformations in conjunction with rotation and stretching transformations. This involves adding a third dimension to centroid $\mathbf{c}_j^{proj}$ and setting this new dimension equal to 1. $\mathbf{r}_j^h$ refers to the homogeneous coordinates along a circle that $\mathbf{c}_j^{proj}$ will be mapped to. For example, if $\mathbf{c}_1^{proj,h}$ refers to observations corresponding to a wind coming at 1 m/s from $\theta = 0$ degrees, then $\mathbf{r}_1^h$ could equal $\mathbf{r}_1^h = [1, 0, 1]^T$; if $\mathbf{c}_2^{proj,h}$ refers to observations corresponding to a wind coming at 1 m/s from $\theta = 90$ degrees, then $\mathbf{r}_2^h$ could equal $\mathbf{r}_2^h = [0, 1, 1]^T$. Solving this problem by taking the gradient of the loss function with respect to $\mathbf{T}$ and setting it equal to 0 yields the minimum mean square error solution in Equation 4.

$$\begin{aligned}
\mathbf{T} &= \mathbf{R}^h \mathbf{C}^{proj,h^T} (\mathbf{C}^{proj,h} \mathbf{C}^{proj,h^T})^{-1} \\
\mathbf{C}^{proj,h} &= \begin{bmatrix} \mathbf{c}_1^{proj,h} & \mathbf{c}_2^{proj,h} & \dots & \mathbf{c}_d^{proj,h} \end{bmatrix} \\
\mathbf{R}^h &= \begin{bmatrix} \mathbf{r}_1^h & \mathbf{r}_2^h & \dots & \mathbf{r}_d^h \end{bmatrix}
\end{aligned} \tag{4}$$

Figure 4d shows our observations $\mathbf{x}_i$ after applying both the projection, $\mathbf{Q}$, and transformation, $\mathbf{T}$. When a new observation comes in, we can apply both $\mathbf{Q}$ and $\mathbf{T}$. Converting the result into polar coordinates yields the azimuth θ of the airflow vector. The magnitude of the vector, r, will then be directly proportional to the wind azimuth magnitude of the wind speed, $|\vec{V}_1|$. To obtain the actual wind speed, we create a simple linear regression model that maps r, observed from the projected and transformed coordinate system, to an actual wind speed value. With both θ and $|\vec{V}_1|$, we have measured the full airflow vector along the azimuthal plane.

Estimating Polar Angle and Wind Speed: Next, we need to find the polar air vector, $\vec{V}_2$, to obtain the full air vector $\vec{V}$ (Figure 2). Figure 5-left shows a scatter plot of the average motor speed (control signal sent from the flight controller) versus polar angle, ϕ, with a constant wind speed, where we see a linear relationship. When there is no wind or when the wind vector is almost entirely in the azimuthal plane, the average control signal sent to the motors centers around a nominal value. When the angle increases ($\phi > 0$), there is a component of wind pushing down on the drone from above. This external force causes the flight controller to actuate motors faster and generate more lift to counteract this downward force. When the angle decreases ($\phi < 0$), there is a component of wind pushing upward on the drone, which helps lift the drone further up into the air. As a result, the flight controller decreases motor speeds to prevent the drone from moving higher. To obtain the airflow component along the z-axis, $\vec{V}_2$, we create a mapping between average motor speed control signals sent from the flight controller by fitting a linear regression curve to the wind speed along the z-axis to obtain $\vec{V}_1$. After estimating $\vec{V}_1$ and $\vec{V}_2$, we can obtain the final 3D airflow vector $\vec{V}_{wind} = \vec{V}_1 + \vec{V}_2$.

Selection of Drone: Existing works generally use large drones and measure wind speeds based on the wind triangle principal. Like the wind triangle method, Anemoi requires the wind to cause noticeable perturbations to the drone to obtain an accurate reading. Because airflows indoors tend to be much smaller than outdoor airflows, a smaller drone would experience greater perturbations and would be more suitable for our application. Moreover, smaller drones are, in general, more suitable for indoor applications because they can more easily fit through narrower indoor spaces and are less likely to damage the environment or cause bodily harm if the drone loses control. As such, we implement Anemoi

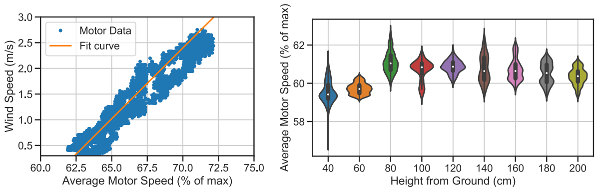

Figure 5: Average motor speed of all four motors versus the wind speed along the z-axis. We see that there is a linear relationship that we can use to obtain wind vector in the z-axis, $\vec{V}_2$ (left) and average motor speed (control signal) at different heights with no wind (right).

on the cheap (< 100 USD), small (30g), and open-source ATK-MiniFly quadcopter [2] shown in Figure 2.

Height of Drone: One concern that arises when we decide to use the average speed of each motor to estimate the Z component of the wind vector is the height of the drone. Intuitively, the higher the drone flies, the faster the propellers need to spin to generate enough thrust to maintain its hover. The typical height of a modern residential room is approximately nine feet tall. Figure 5-right shows the average speed of the motors as we vary the height of the drone between 1 to 8 feet (30-210cm) from the ground, which is at least 1 foot away from both the ceiling and ground to ensure that the drone can maintain a stable hover. We see that the average motor speed in this low-height range is relatively stable, and thus we do not need to heavily account for the height in typical indoor scenarios.

Effects of Propellers on Estimation: Because the propellers generate an airflow to propel the drone in the air, a concern is that this self-generated airflow will affect measurements of external airflows. The drone's PID controller is tuned to incorporate the effects of its own motors/propellers and any corrections it applies are a reaction to external effects, such as external wind. As such, the wind generated by the propellers does not affect Anemoi's wind speed measurements, which is consistent with our own observations and evaluation (Section 5).

Advantages over Traditional Drone Sensing: Another solution to autonomously mapping airflows throughout a 3D space is to attach an additional wind sensor on the drone. There are two major drawbacks to this approach. First, in constrast to Anemoi, the wind generated by the propellers would significantly impact the external sensor's measurement. Second, most airflow sensors and anemometers only provide wind speed and lack direction estimation. Airflow sensors that can provide direction are generally very large and not suitable for drones to carry [16]. For example, the Minifly drone and many other MAVs are designed to be extremely light-weight, only carrying a payload of several grams (<15g), making it difficult to attach an external sensor while keeping the drone stable. Because of these reasons, we do not compare Anemoi against this method in Sections 5 and 6; rather, we compare our sensorless method against solutions that need to leverage sensors typically found on drones (e.g., optical flow and IMU).

4 AUTONOMOUS ESTIMATION OF AIRFLOW FIELDS

Since airflow varies continuously over space and time, it is impossible to measure at every possible location all the time. Therefore, we need a strategy to extrapolate airflow in a target volume as accurately as possible, with a limited number of samples, in the shortest amount of time. In this section, we propose an information theoretic approach to sampling and leverage *computational fluid dynamics* (CFD) to perform extrapolation.

To accurately measure the airflow field, it is important to efficiently sample the space in a limited amount of time. While CFD is a well-studied field with accurate methods to extrapolate fluid flow [53, 73], it is computationally expensive to use in practice. However, in order to select the next location to sample, we only need the effect of that sample on the overall estimated uncertainty (entropy) of the full field. Therefore, we decouple this challenge into two tasks - (i) real-time entropy estimation for exploration (Section 4.1) and (ii) post-processing that computes the overall airflow field with CFD (Section 4.2).

4.1 Information Gain based Exploration

Exploration or *Active Mapping* [11, 15, 63] in robotics is the problem of driving the robot to a set of locations so as to maximize the estimate of a function of interest (airflow, in our case) in the target space. With each measurement, the drone gains information about the air velocities in the target space, and improves the system's capacity to better estimate the airflow field in the entire space, while reducing entropy or uncertainty. Entropy is the measure of expected information that the value of x, from a probability distribution $p(x)$, carries for a random variable X [70], given by:

$$H_p(x) = E[-\log p(x)] = -\int p(x) \log p(x) \, dx \tag{5}$$

For a Gaussian distribution, entropy is directly proportional to its variance. In other words, the entropy of noisier and less certain estimates are higher. We discretize our space into a 3-D grid with a configurable resolution and make the following assumptions: i) the airflow of all points in a single cell are the same, ii) the entropy of different cells are independent, iii) the entropy of a cell is *maximally reduced* if Anemoi measures airflow in that cell.

We propose the following information-theoretic approach for path planning. Anemoi selects the cell, c_j, with the highest *utility* as the next cell to measure airflow, as shown in Equations 6, 7, and 8.

$$\text{Utility}_{c_j} = \text{Information Gain}_{c_j} - \text{Cost}_{c_j} \tag{6}$$

$$\text{Information Gain}_{c_j} = \sum_{i=1}^{N_{c_j}} H_p(c_i) \tag{7}$$

$$\text{Cost} = \|X_{c_j} - X_{\text{drone}}\| \tag{8}$$

where the information gain of cell c_j is the sum of entropies in the N_{c_j} neighboring cells, and the cost is the euclidean distance the drone needs to travel between its current position, X_{drone} to arrive at the location of cell c_j, X_{c_j}. In other words Anemoi selects the cell that has the best balance between a high information gain (or uncertainty) and a low cost. We model the cost as the distance Anemoi needs to travel because the drone has a limited flight time that diminishes the further it needs to travel.

Once Anemoi samples a cell, c_m, its entropy is maximally reduced. However, the airflow observed in cell c_m also gives insight into the airflow in neighboring cells, c_i. As such, we propose the following formulation entropy reduction in neighboring cells after each measurement:

$$H_p(c_i) = H_p(c_i) \ - \ |S(F_m, C_m - C_i)| \ e^{\frac{-\|C_m - C_i\|}{\lambda \|F_m\|}} \tag{9}$$

$$S(A, B) = \frac{A.B}{\|A\| \ \|B\|} \tag{10}$$

where F_k is airflow velocity at cell k, c_k is cell k, C_k is the coordinates of cell k, and λ is the attenuation factor.

In this formulation, $|S(F_m, C_m - C_i)|$ incorporates the insight that the measurement at c_m provides more information to cells that are aligned in the direction of measured airflow, F_m; hence, the entropy of cells that are physically in the direction of the moving airflow are reduced more than those in other directions. $\|F_m\|$ scales the reduction with the magnitude of the air velocity and $\|C_m - C_i\|$ scales the entropy reduction inversely with distance from the measured cell, accounting for noisier estimates at distance.

4.2 Computational Fluid Dynamics for Extrapolating Airflow Field

After obtaining all measurements, we need to extrapolate airflow measurements to cells that are not measured to obtain the airflow field across the entire space. To accomplish this, we leverage computational fluid dynamics (CFD), which is the field of analytically or numerically modeling natural phenomena such as airflow, temperature, and pressure [43].

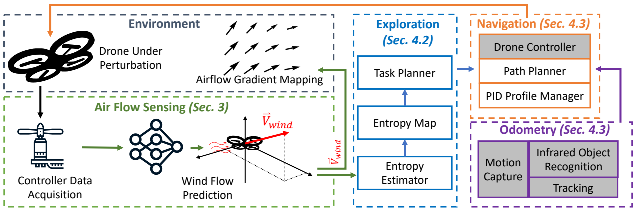

Figure 6: System architecture of Anemoi. Components shaded in gray are from existing literature.

We model the airflow throughout our space leveraging the *Viscous Burger equation* [12], a simplified version of the set of partial differential equations known as the *Navier-Stokes equations* [41] that are commonly used to model the flow of viscous fluids. Specifically, we leverage the stochastic version of the Burger's equation, for which the 1-D version of the equations are shown here:

$$u_i^{t+1} = u_i^t - u_i^t \frac{\Delta t}{\Delta x}(u_i^t - u_{i-1}^t) + \nu \frac{\Delta t}{\Delta x^2}(u_{i+1}^t - 2u_i^t + u_{i-1}^t) \tag{11}$$

$$\frac{\partial u}{\partial t} + u \frac{\partial u}{\partial x} = \upsilon \frac{\partial^2 u}{\partial x^2} - \lambda \frac{\partial \eta}{\partial x} \tag{12}$$

This equation shows the change in the u component of airflow velocity w.r.t to time t and the spatial dimension x. With a known initial value u^0, in an area, the velocity vector can be extrapolated over discrete time steps Δt with a viscosity constraint υ over the spatial step Δx. *The initial values, u^0, are directly measured by Anemoi* as the drone traverses the space (Section 4.1). Additionally, this formulation accounts for the noise of extrapolated fields, $\eta(x, t) = \dot{W}(x, t)$, modeled as a Wiener process (Brownian noise) [76], where estimates of airflow further from a measured location are less certain. Through these methods, we estimate the entire airflow field from discrete measurements of airflow vectors. In our evaluation (Section 6), we run this extrapolation after every airflow vector measurement to assess accuracy of the estimated airflow field at each iteration.

4.3 Realizing Robot Exploration on Palmtop UAVs

Anemoi requires accurate and real-time odometry to safely navigate the drone to new locations. Further, we need a set of trajectory commands that navigate the drone through the planned path. Finally, because the drone is relatively light, to keep it navigating safely in hazardous winds, we switch between a navigation PID profile that is less aggressive and a measurement PID profile that is desirable for the prediction model. Our overall software pipeline is shown in Figure 6.

Odometry: To isolate sensing error, we rely on external sensing from a Vicon motion capture system. This system provides accurate real-time odometry. This information is sent to the drone via radio, and fused with its onboard state estimator. There is a large body of work in obtaining odometry from cameras and LiDARs [5, 7, 8, 19] onboard. We characterize optical-flow based odometry error to demonstrate its feasibility in Section 6. Integrating Visual SLAM [8] was not feasible on the current system due to interference of Wi-Fi packets from the camera module and the radio signals

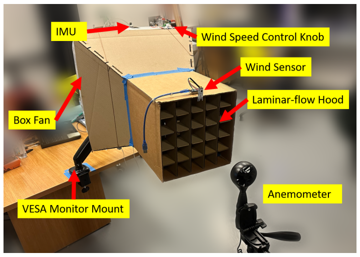

Figure 7: Experimental setup for generating stable (laminar) airflows at various velocities from the drone.

for control. However, we envision that such integration in a comparable micro-aerial vehicle is possible.

Path Planning: We have modified onboard firmware on the drone to obtain global location in real-time from the external camera system. Given the drone's current location and its destination, we generate a trajectory by segmenting the goal path into 10 cm pieces. This method can also perform obstacle avoidance using the external sensing system. The intermediate way-points are published at a fixed interval of 750ms to the drone's state estimator.

Controller Profile Manager: Because of the drone's weight, we observed empirically that a less aggressive feedback controller helps the drone stably navigate in a windy environment. However, to measure airflow, our feedback controller needed to be aggressive. To manage the two profiles, we created an onboard profile manager, which automatically switches between navigation and measurement profiles. The manager determines that the drone is in position and stable, when position error vector norm and velocity vector norm are both under some threshold.

5 EVALUATING AIRFLOW VECTOR ESTIMATION

In this section, we describe our experimental setup and evaluate how well Anemoi estimates airflow vector at a location in a controlled environment. We evaluate the performance of autonomous sampling and extrapolation in Section 6.

5.1 Experimental Setup

To compare how well Anemoi estimates airflow vectors in various configurations, we need to generate stable airflows. We used a 20 by 20 inch box fan mounted on a computer monitor stand with an inertial measurement unit (IMU) to obtain the direction of the airflow. However, there are several challenges that need to be addressed to conduct a thorough evaluation. Our full experimental setup is shown in Figure 7.

5.1.1 Generating Stable Flows with Fine Control.

5.1.1 Generating Stable Flows with Fine Control. Most flows found in the real world are turbulent, meaning that they are chaotic due to rapidly changing air velocities and pressures. This makes it difficult to obtain ground truth air velocity measurements. To address this challenge, we built a laminar hood based on [20]. We observe that adding the hood significantly reduced the turbulence in the wind; without the hood, the wind speed fluctuated with a variance of 0.2m/s, compared to 1.1m/s while using the hood. These measurements were taken 1.5 feet away from the fan/hood opening.

Additionally, most indoor and home fans are adjustable to only around three different speeds, whereas we need fine-grained control to evaluate estimation errors across a wide range of speeds. As such, we outfit our fans with an adjustable voltage regulator to digitally control the wind speed.

5.1.2 Ground Truth Wind Speeds. In addition to generating stable winds, we need to measure the ground truth wind vector at the location of the drone. However, the MiniFly drone is too light and cannot carry an additional wind sensor to measure ground-truth wind speeds in real-time. We can place a wind sensor at the mouth of the hood, but we need to measure the wind speed at the location where the drone will be hovering. Therefore, we use an anemometer placed 1.5 meters from the opening of the hood to measure ground truth wind speeds in addition to the wind sensor placed at the mouth of the hood, as shown in Figure 7, and create a model that maps readings from the wind sensor at the mouth of the hood to the ground truth wind speed measured at the anemometer. When collecting data with Anemoi, we remove the anemometer and hover the drone in its place while changing wind speeds and carefully moving the box fan to different directions, while maintaining the same distance.

By carefully moving the box fan to different directions, while maintaining the same distance from the location of the anemometer, we can ensure consistent wind speeds at the same box fan setting. Moreover, using the anemometer to create a model that maps the wind speed at the mouth of the box fan hood to the location of the anemometer allows us to have an accurate wind speed reading even when the anemometer is replaced by the drone.

5.2 Effects of Tuning PID

Anemoi measures airflow based on how the drone actuates the motors in response to perturbations caused by the wind, which can be influenced by selecting different PID gain constants. Figure 8 - top shows the default architecture of the PID controller for the MiniFly drone, which we denote as cascade control (CC) in Figure 8a. The inertial and body frame both leverage two cascaded controllers: controllers for correcting position and velocity (inertial frame) and controllers for correcting attitude (roll, pitch, yaw) and angular rate (body frame). We found very little change in performance when we adjusted the PID constants for the position and velocity control; this is because the drone attempts to hold its own position while taking measurements, so the position and velocity require fewer corrections compared to the attitude of

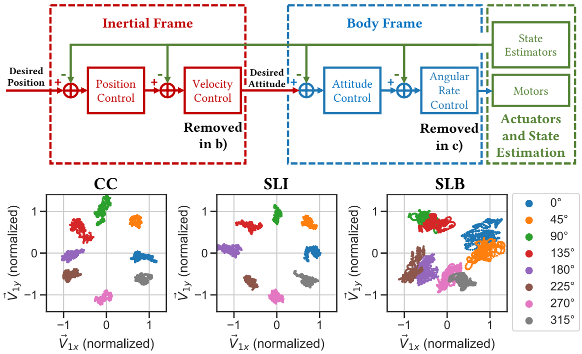

(a) (b) (c)

Figure 8: Top) PID control schemes evaluated. Azimuth wind velocity mapping for a) Cascade control (CC). b) Single loop inertial frame (SLI). c) Single loop body frame (SLB).

the drone. Since the angular rate has the highest impact on the stability of the drone, which could greatly impact factors such as sensitivity and accuracy of measurements, we study the effects of PID constants for angular rate.

Figure 9a shows a scatter plot of estimated azimuth directions and wind speeds with varying angular rate PID constants. We first vary the P constant, while removing I and D, (Figure 9a - c) and see that as we increase P, the angle estimation improves and becomes more separable, but the speed estimation sees higher error, as shown by higher variance of the points spreading out from the center of the plot. This is because increasing the P constant causes the controller to compensate for errors more aggressively, causing the controller to adjust motor speeds more frequently and with greater variations. We set $P = 3000$ because it gives the best trade-off in angle and wind speed estimation.

Next, we tune D and see a similar trend that increasing D decreases wind speed estimation. However, we also see that increasing D decreases direction estimation accuracy and we observe that the drone becomes much more unstable very quickly as it attempts to anticipate instabilities from future wind perturbations. As such, we select the lowest D possible that keeps the drone relatively stable, $D = 65$.

Finally, we vary I and see very little impact on speed and direction error compared to increasing D. This is because I corrects for past errors or biases in the angular rate of the body frame. However, our drone hovers in place while taking a measurement and must keep a stable attitude to maintain its hover. Because it is able to do this under most typical wind speeds in indoor environments even without high levels of I, increasing I has relatively low impact on our measurements.

5.3 Effects of Different PID Control

One of our aims is to show that Anemoi can be adapted into most commercial off-the-shelf micro-air vehicles without making any modifications. There are a number of open source and commercial drones that leverage the same or similar type of cascade and nested PID control scheme as the MiniFly, suggesting that Anemoi can be easily adapted into many existing drones. For example, the Crazyflie drone platform has almost the exact same dimensions/weight and leverages the same PID control scheme for correcting errors in position and attitude [25]. PX4 [68] and ArduPilot [67], two of the most common open source flight controllers for UAVs, as well as platforms that leverage these flight controllers (e.g., [4, 22, 45]) also leverage the same cascade and nested control structure. Many other custom drone platforms still leverage or have the option to use similar schemes [38, 66].

We implement two additional control schemes based off of MiniFly's cascade control (Figure 8) for comparison by removing the velocity controller (single loop inertial frame or SLI) and removing the angular rate controller (single loop body frame or SLB). Figure 8a-c shows the scatter plot of estimated azimuth direction and wind speeds comparing cascade control (CC) with SLI and SLB, respectively. We see that when the velocity controller is removed (SLI), there is little change in performance because the position of the drone changes more slowly than the attitude, which has more impact on sensing performance. However, when the angle rate controller is removed (SLB), the variance of the estimated flows increase because the drone becomes less stable and reacts more to smaller wind speeds. In the following subsections, we compare all three control schemes in terms of the sensitivity and range of wind speeds they can measure.

5.4 Range and Sensitivity

Wind speeds that are either not strong enough to cause the drone's PID controller to noticeably account for perturbations or too strong that it causes the drone to completely lose control will result in high errors. In this section, we analyze two aspects of Anemoi to try and maximize the range and sensitivity of speeds we can detect.

We want to find the *minimum measurable speed* and the *maximum measurable speed* that Anemoi can detect. To maximize this range, we tune the PID constants to influence how the drone accounts for perturbations from airflows.

Figure 10 plots a heat map of the wind speed error of Anemoi versus the ground truth wind speed, while varying the control scheme and the PID gain constants for angular rate control in the CC (balanced, high I, and high D constants). We use a window size of 1000ms, which we empirically determine in Section 5.5. The typical maximum speed coming out of a fan or HVAC unit can be up to 4 – 5m/s directly next to the unit [28]. However, a drone would not typically fly this close to a wall or object. As such, we measure and plot wind speed errors up to 3m/s as this was the maximum wind speed of several different AC and fan units we measured at around 1 meter away.

We see that many curves follow a similar pattern: below and above certain ground truth wind speed, s_{min} (start) and

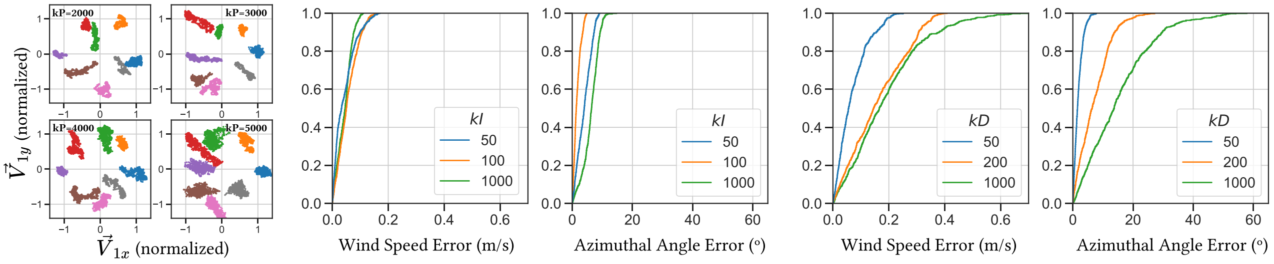

(a) kP **(b) kI** **(c) kD**

Figure 9: a) Scatter plots of estimated wind speed and direction varying the P gain constant. b) Left - CDF of wind speed error, Right - CDF of azimuth angle error varying the I gain. c) Left - CDF of wind speed error, Right - CDF of azimuth angle error varying the D gain.

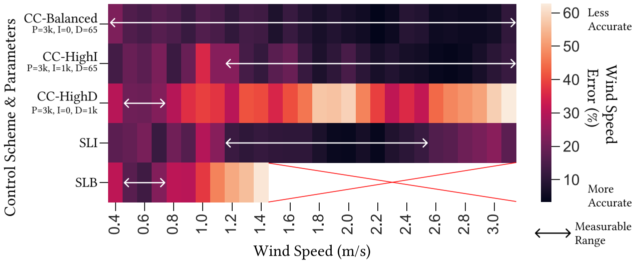

Figure 10: Estimated wind speed error versus ground truth at varying PID values and control schemes.

s_{max} (end), the wind speed estimation error drastically increases. This is because the wind is either too weak to perturb the drone significantly ($< s_{min}$), or too strong that it blows the drone away ($> s_{max}$). Below 0.4m/s our ground truth anemometer was not sensitive enough to measure. Additionally, several configurations (e.g., CC-Balanced and CC-HighI) have an upper wind speed, s_{max} above 3m/s, which is above the typical wind speeds in an indoor environment.

We define s_{max} as the *maximum measurable speed*, where all wind speeds greater than s_{max} yield an error greater than 1.5 times the *lowest measured error*, l_{err}, achieved. We define 1.5 times the *lowest measured error* as the *maximum stable error*, where wind speeds that yield errors less than this value are relatively accurate and stable; we determine the 1.5 multiplicative factor empirically based on looking at the heat map of each method's wind speed error versus ground truth wind speed in Figure 10. The *minimum measurable speed*, s_{min}, is the speed where all wind speeds less than s_{min} show this same behavior. In Figure 10, we highlight the *measurable range* or the region between the maximum and minimum measurable speeds.

When comparing the effects of I and D on the estimation range for the CC scheme, we see that a high D changes the sensing characteristics the most. The drone is more sensitive to lower wind speeds as it becomes less stable with less external force. However, the accuracy for higher wind speeds reduces. Changing to the SLI control scheme, we see that the performance is more similar to CC-Balanced and CC-HighI as we observed that the velocity controller (removed in SLI) has a small effect on the stability of the drone while it is hovering in place. However, removing the angular rate controller (SLB) reduces the stability of the drone more than any other configuration. As such, it is much more sensitive to lower wind speeds, but cannot even maintain a hover at higher wind speeds (denoted by the "X").

We select the CC-Balanced configuration ($P = 3000$, $I = 0$, and $D = 65$) because it covers the highest range of wind speeds compared to other configurations and because it has the best trade-off in wind speed and direction estimation accuracy, as discussed in Section 5.2.

5.5 Estimation Error, Window Size, and Latency

There are many parameters that affect the overall accuracy and performance of wind vector estimation, including control signals to use, features to extract, mapping models, and measurement window sizes. We evaluate them below.

Estimation Method: Figure 11 shows 25th, 50th, and 75th percentile error for estimated wind speed, azimuthal angle, and polar angle for Anemoi compared against a state-of-art method for estimating airflow using a drone without a wind sensor, which we denote as BL [46]. This method models the aerodynamics of the drone and utilizes the wind triangle method to obtain wind speeds by estimating the drone's air velocity using an IMU and the ground velocity using GPS. Because GPS is highly inaccurate indoors, we leverage an optical flow to obtain ground velocity based on [58]. This method cannot obtain measurements along the polar axis. Overall, Anemoi achieves 0.11m/s wind speed, 6.19° azimuth angle, 8.40° polar angle average estimation errors, which is 0.41m/s and 25.1° lower than the state-of-the-art azimuth velocity error.

We estimate the error by varying the wind speed between each configuration's *maximum measurable speed* and *minimum measurable speed*, which defines the *range* of wind speeds that the system can sense. Beyond these speeds, the

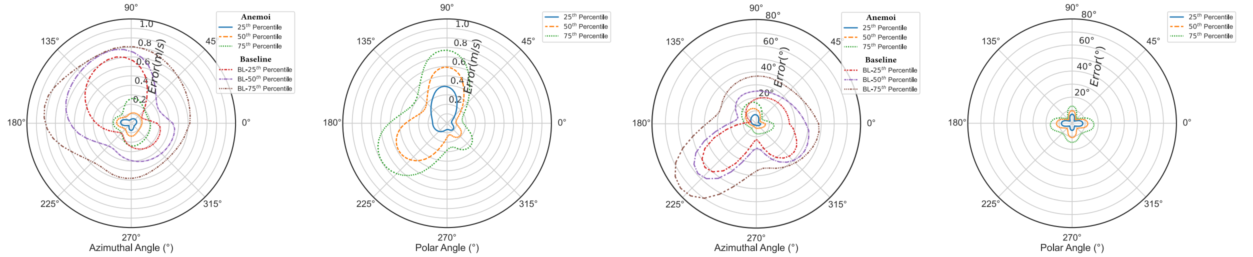

(a) Wind speed error vs. azimuth angle of wind. **(b) Wind speed error vs. polar angle of wind.** **(c) Azimuthal angle error.** **(d) Polar angle error.**

Figure 11: Error of estimated wind speed vs. the direction of incident wind with varying azimuth (a) and polar angle (b). Error of estimated azimuth angle of wind varying the azimuth of incident wind (c) and error of estimated polar angle of wind varying the polar angle of incident wind (d). Here we set P = 3000, I = 0, D = 65, window size = 1000ms, which is the configuration we used to perform autonomous airflow mapping (Section 6).

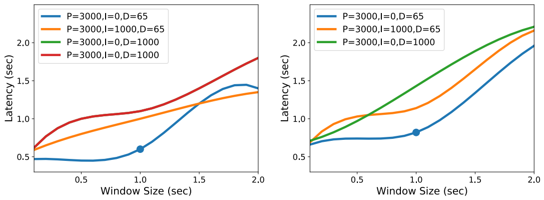

Figure 12: *Latency* of wind speed measurements for varying window sizes (left) and *latency* of wind angle measurements with varying window sizes (right).

error drastically increases because the wind speed is either too low to significantly affect the drone or too strong that the drone loses control. We make this comparison using the following parameters for Anemoi: P = 3000, I = 0, D = 65, window size = 1000ms, which we also use to perform autonomous airflow mapping.

Window Size and Latency: How many samples we choose to average together may also impact the quality of measurements. For example, using only 1ms worth of acceleration data may not be enough time to capture the dynamics of the drone to give an accurate reading. However, using an extremely long window (e.g., 10 seconds) increases the time it takes for the sensor to adjust to a new wind speed/direction.

Figure 12 plots the *latency* for Anemoi to begin measuring wind speeds with low errors in the *measurable range* (defined in Section 5.4), from when the wind speed or direction changes vs. the window size taken for each measurement.

Latency refers to the amount of time from when we turn on the fan to when Anemoi reads a wind speed less than the *maximum stable error* (defined in Section 5.4) and/or when Anemoi begins to read the correct wind direction (azimuthal and polar) within +/-10 degrees.

We select a window size of 1000ms (marked as a blue dot), which yields a latency of around 0.6 seconds for wind speed and 0.8 seconds for wind direction (varying both azimuth and polar) for the configuration we use to perform autonomous mapping ($P = 3000$, $I = 0$, and $D = 65$). Most scenarios will involve the drone taking measurements for several seconds at one location before moving to the next location, so a sub-1 second latency is acceptable.

6 EVALUATING AIRFLOW FIELD ESTIMATION

In this section, we evaluate the information theoretic autonomous sampling method and CFD-based airflow field extrapolation. We use simulations to isolate sampling and extrapolation performance from airflow estimation errors. We use real-world experiments to make end-to-end evaluation that include the performance of single point estimates, sampling and extrapolation. For both simulation and real-world scenarios, we divide the space into 1mx1mx1m grids for exploration.

For the simulations, we generated ground truth airflow in a 10mx10mx10m space using CFD with fixed airflow sources. We applied the information gain-based exploration algorithm to sample the space. We then used these point measurements to estimate the airflow field in the space with CFD. We compare our approach against a random sampler with the same number of samples. We note that a random sampler is a better baseline than predetermined routes like lawn-mower pattern for applications that extrapolate signals with spatial variations with limited samples [27].

Real world airflow experiments are hard to construct, control and measure ground truth. For these experiments, we used two sources of air intake in a 3mx3mx3m grid. We took great care to ensure parity in the ground truth and drone measurements and to ensure that no additional wind intakes or outlets affected the measurements. Further, we isolated the localization error from the sensing error by using a motion capture system for position estimates.

To quantify the performance of our exploration, we use three metrics: (i) **Entropy Per Cell**: Total entropy in the system averaged over the number of cells in the target area.

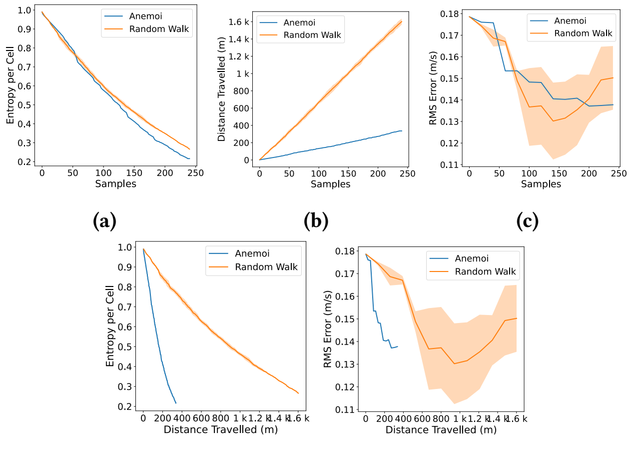

Figure 13: Performance of our entropy-based exploration in simulation vs Random walk. (a) Average entropy per cell after each step (lower is better). (b) Cumulative distance travelled. (c) RMS Error of estimated airflow (d) Average entropy per cell for distance traveled. (e) Error reduction of estimated airflow for distance traveled (lower is better).

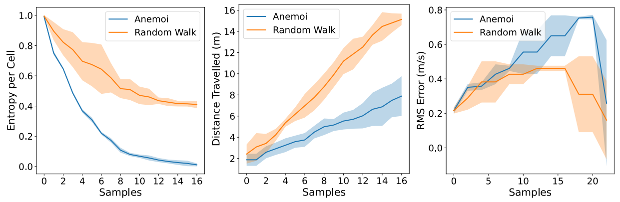

Figure 14: Entropy-based exploration experiments - Average Entropy per cell per step (lower is better)(left), Cumulative distance travelled (center) and RMS Error of estimated airflow (right).

It is an indicator of the amount of information gained by the sampling process. A valid exploration algorithm must reduce the entropy per cell after each step of the exploration process. A lower entropy per cell indicates a better choice of sampling position. Further, this metric should be reduced by traveling the minimum distance, (ii) **Distance Travelled**: The total distance travelled by the drone during the measurement process is a good indicator of the cost of making measurements. A lower distance travelled would result in a faster execution with fewer stops for recharging, and (iii) **RMS Error in Air-flow estimation**: Air-flow can be estimated with any number of initial samples with CFD. For a given number of samples, a higher sample quality would result in a better estimation of air-flow in the target space.

We observe from Figure 13 that both the exploration-based sampling and random sampling reduce the entropy in the system for every sample. However, we note that the cost of sampling in terms of distance travelled increases much more rapidly in random sampling. Our exploration method considers the cost of sampling in the decision-making process, performing better for a given number of samples (Figure 13a) and the same distance traveled (Figure 13d). We also observe from Figure 13c that the RMS error computed with the ground truth and estimated airflow reduces predictably with exploration, whereas for random sampling it is not as predictable. This suggests that the quality of samples from exploration is better. Finally, the total distance traveled by the exploration algorithm is far lesser than random to achieve the same RMS error reduction in the estimate (Figure 13e).

Figure 14 shows Anemoi's performance in real-world experiments. We see that the performance trends for entropy reduction and cost are similar compared to simulations. Given the limited information about the overall airflow at the start, our method helps rapidly decrease the entropy per cell. Our exploration algorithm is able to intelligently pick sampling locations balancing information gain and distance cost, minimizing total distance traveled. We find that due to the turbulent nature of wind in the real world, airflow measurement is noisy in both ground truth and Anemoi measurements. We observe that the general trend of the airflow is captured as shown in Figure 15. However, we observe that as additional samples are collected, our overall RMS error increases (Figure 14). On close examination, we attribute this to errors in the direction of sampled airflow, particularly in points where the measured airflow is very low. This causes our CFD to compute overall airflow that deviates from the ground truth. In the future, we can use our calibration from Figure 10 to filter out values outside our sensor sensitivity for improved airflow estimation.

7 DISCUSSION

Sensor Sensitivity: One challenge in designing the airflow estimation on the drone was the range of values it was able to sense. Our palmtop drone was sensitive to higher volume of airflow, such as when it is in front of a box fan. It was not very sensitive at greater than 6-8m away from an airflow source. We suspect this is an artifact on the drone size. A future research direction we hope to pursue is the study of drones of various sizes and the limits of their sensitivity.

Different Control Schemes: In this work, we explored how different PID control schemes affect Anemoi. Although PID control schemes are most commonly used, there are a variety of drastically different drone control schemes that have been proposed previously, such as reactive control [9] and INDI [66], which we plan to explore in future work.

Exploration Sensitivity: To compute in real-time, the exploration algorithm chooses sensor poses based on maximal entropy reduction in unmeasured locations. The actual entropy reduction after measurement depends on the magnitude and direction of the airflow at the chosen location. This

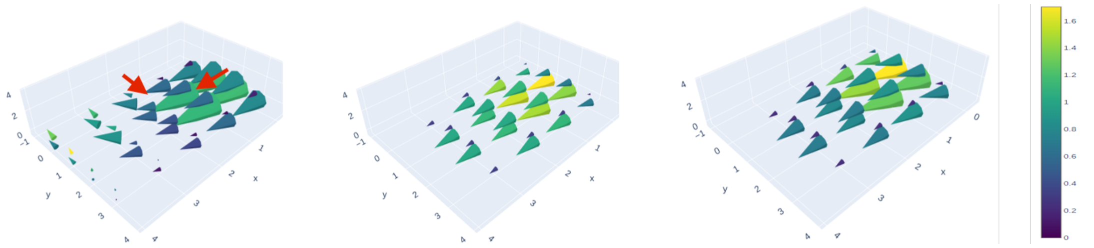

Figure 15: Airflow from real-world experiments. Ground truth airflow generated with two sources indicated with red arrows (left), Airflow estimated by running CFD on measurements made with the Anemoi exploration algorithm (middle), and Airflow estimated by running CFD on measurements made with random strategy (right).

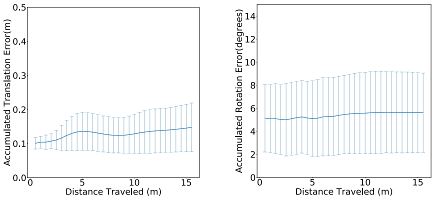

Figure 16: Optical-flow based localization error across five flight trials - Accumulated Translation Error (left) and Accumulated Rotation Error (right).

is a trade-off to achieve online execution. We assume that airflow in the entirety of the target space is estimated as a post-processing step after sampling, since airflow extrapolation using CFD is a time-consuming process.

Full autonomy: While complete onboard autonomy was desirable on the drone during the sampling phase, we relied on external sensing for odometry. The primary reason is the limited battery capacity on small drones for onboard odometry. Our empirical experience suggests that we might need to move to a larger drone with a bigger battery to incorporate full autonomy in airflow estimation.

Localization System: While we used a motion capture system to isolate localization error from sensing error, an onboard localization system can be used in practice [8, 33, 79]. In Figure 16, we can see the cumulative optical-flow and IMU based localization errors on our drone is very low while traveling our experimental setup (< 0.3m). Since we divided the space into $1m^3$ grids, the average localization error should be less than $0.5m$ to ensure that the drone stays within the grid while measuring; a drone that drifts outside of the grid during measurement will be affected by airflows from neighboring grids, which reduces measurement accuracy. From Figure 16, we see that the localization error rises during operation but does not go above $0.3m$. One aspect that could be used to alleviate this problem is to take advantage of the drone's limited flight time. When it returns to a base station to recharge or swap batteries, the drone could be reset, which would also reset sensor drifts and reduce localization errors.

Applications: Our platform has been evaluated for and can be used to monitor health of indoor spaces by measuring ventilation and air flow. A $1m$ grid is suitable for such applications [13, 61] where predominantly CFD analysis is done over measurements of sources and vents or through manual measurements in fixed points $> 1m$ apart. In contrast, finer grained measurements of up to $0.5m$ grids may be required for other applications like air flow based thermal comfort analysis [54]. For such applications, a higher localization accuracy is required.

8 CONCLUSION

We present Anemoi, a low-cost sensorless indoor drone system for autonomous estimation of 3D airflow fields. We demonstrate that Anemoi is able to estimate airflow without the use of external sensors by training a data-informed geometric model that maps motor control signals (under perturbation) to known (generated) wind vectors. We show the effects of system parameters, such as PID values, control signals/features to use, mapping models, and measurement window sizes, on the performance of the system, including range, sensitivity, estimated error, and latency. In comparison with a baseline state-of-art drone-based method, Anemoi achieves up to 0.41m/s and 25.1° average lower errors in wind speed and angle estimation. We propose an exploration algorithm that selects optimal points in the space for measurements, and demonstrate that Anemoi is able to map the 3D airflow field with an average RMS error of 0.73m/s and average angular error of 29.79°.

ACKNOWLEDGMENTS

This research was partially funded by the National Science Foundation under Grant Numbers CNS-1943396 and CNS-1846320, as well as COGNISENSE, one of seven centers in JUMP 2.0, a Semiconductor Research Corporation (SRC) program sponsored by DARPA. The views and conclusions contained here are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of Columbia University, University at Buffalo, NSF, SRC, DARPA, or the U.S. Government or any of its agencies.

REFERENCES

[1] Andrea Alaimo, Valeria Artale, Cristina Lucia Rosa Milazzo, and Angela Ricciardello. 2014. PID controller applied to hexacopter flight. *Journal of Intelligent & Robotic Systems* 73 (2014), 261–270.

[2] Alientek. 2022. Punctual atomic MiniFly quadcopter open source UAV flight control STM32 DIY combat kit. https://m.tb.cn/h.U8v3EdX?tk=3F7fdWER5DH.

[3] V Madhu Babu, Kaushik Das, and Swagat Kumar. 2017. Designing of self tuning PID controller for AR drone quadrotor. In *2017 18th international conference on advanced robotics (ICAR)*. IEEE, 167–172.

[4] Tomas Baca, Matej Petrlik, Matous Vrba, Vojtech Spurny, Robert Penicka, Daniel Hert, and Martin Saska. 2021. The MRS UAV system: Pushing the frontiers of reproducible research, real-world deployment, and education with autonomous unmanned aerial vehicles. *Journal of Intelligent & Robotic Systems* 102, 1 (2021), 26.

[5] T. Bailey and H. Durrant-Whyte. 2006. Simultaneous localization and mapping (SLAM): part II. *IEEE Robotics Automation Magazine* 13, 3 (Sep. 2006), 108–117. https://doi.org/10.1109/MRA.2006.1678144

[6] RJ Barthelmie, L Folkerts, FT Ormel, P Sanderhoff, PJ Eecen, O Stobbe, and NM Nielsen. 2003. Offshore wind turbine wakes measured by sodar. *Journal of Atmospheric and Oceanic Technology* 20, 4 (2003), 466–477.

[7] Ali J. Ben Ali, Zakieh Sadat Hashemifar, and Karthik Dantu. 2020. Edge-SLAM: Edge-Assisted Visual Simultaneous Localization and Mapping. In *Proceedings of the 18th International Conference on Mobile Systems, Applications, and Services* (Toronto, Ontario, Canada) *(MobiSys '20)*. Association for Computing Machinery, New York, NY, USA, 325–337. https://doi.org/10.1145/3386901.3389033

[8] Ali J. Ben Ali, Marziye Kouroshli, Sofiya Semenova, Zakieh Sadat Hashemifar, Steven Y. Ko, and Karthik Dantu. 2022. Edge-SLAM: Edge-Assisted Visual Simultaneous Localization and Mapping. *ACM Trans. Embed. Comput. Syst.* (aug 2022). https://doi.org/10.1145/3561972 Just Accepted.

[9] Endri Bregu, Nicola Casamassima, Daniel Cantoni, Luca Mottola, and Kamin Whitehouse. 2016. Reactive control of autonomous drones. In *Proceedings of the 14th Annual International Conference on Mobile Systems, Applications, and Services*. 207–219.

[10] KA Browning and R Wexler. 1968. The determination of kinematic properties of a wind field using Doppler radar. *Journal of Applied meteorology and climatology* 7, 1 (1968), 105–113.

[11] Wolfram Burgard, Mark Moors, Cyrill Stachniss, and Frank E Schneider. 2005. Coordinated multi-robot exploration. *IEEE Transactions on robotics* 21, 3 (2005), 376–386.

[12] J.M. Burgers. 1948. A Mathematical Model Illustrating the Theory of Turbulence. Advances in Applied Mechanics, Vol. 1. Elsevier, 171–199. https://doi.org/10.1016/S0065-2156(08)70100-5

[13] John Kaiser Calautit and Ben Richard Hughes. 2014. Measurement and prediction of the indoor airflow in a room ventilated with a commercial wind tower. *Energy and Buildings* 84 (2014), 367–377. https://doi.org/10.1016/j.enbuild.2014.08.015

[14] Xiaodong Cao, Junjie Liu, Nan Jiang, and Qingyan Chen. 2014. Particle image velocimetry measurement of indoor airflow field: A review of the technologies and applications. *Energy and Buildings* 69 (2014), 367–380.

[15] Shengyong Chen, Youfu Li, and Ngai Ming Kwok. 2011. Active vision in robotic systems: A survey of recent developments. *The International Journal of Robotics Research* 30, 11 (2011), 1343–1377.

[16] R. M. Young Company. 2021. 3D Ultrasonic Anemometer. https://www.youngusa.com/product/ultrasonic-anemometer/.

[17] Ery Djunaedy, JLM Hensen, and MGLC Loomans. 2003. Toward external coupling of building energy and airflow modeling programs. *ASHRAE Transactions* 109, 2 (2003), 771–787.

[18] Richard J Doviak, Peter S Ray, Richard G Strauch, and L Jay Miller. 1976. Error estimation in wind fields derived from dual-Doppler radar measurement. *Journal of Applied Meteorology and Climatology* 15, 8 (1976), 868–878.

[19] H. Durrant-Whyte and T. Bailey. 2006. Simultaneous localization and mapping: part I. *IEEE Robotics Automation Magazine* 13, 2 (June 2006), 99–110. https://doi.org/10.1109/MRA.2006.1638022

[20] Pitsco Education. 2016. Laminar Flow with the Air Control Funnel. https://www.youtube.com/watch?v=lgB5KxX3l84. [Online; accessed 09-August-2022].

[21] E+E Elektronik. 2022. Airflow Sensor for Measuring Air Velocity. https://www.epluse.com/products/air-velocity-instrumentation/air-flow-transmitters-and-probes/.

[22] Matthias Faessler, Antonio Franchi, and Davide Scaramuzza. 2017. Differential flatness of quadrotor dynamics subject to rotor drag for accurate tracking of high-speed trajectories. *IEEE Robotics and Automation Letters* 3, 2 (2017), 620–626.

[23] FIXAR. 2020. Drones for Indoor Mapping and Inspection - FIXAR UAV. https://fixar.pro/products/fixar-indoor/.

[24] Y Fujiyoshi, K Yamashita, and C Fujiwara. 2009. Detection of organized airflow in the atmospheric boundary layer and the free atmosphere using a 3D-scanning coherent Doppler lidar. In *International Symposium on Photoelectronic Detection and Imaging 2009: Laser Sensing and Imaging*, Vol. 7382. SPIE, 55–64.

[25] Wojciech Giernacki, Mateusz Skwierczyński, Wojciech Witwicki, Paweł Wroński, and Piotr Kozierski. 2017. Crazyflie 2.0 quadrotor as a platform for research and education in robotics and control engineering. In *2017 22nd International Conference on Methods and Models in Automation and Robotics (MMAR)*. IEEE, 37–42.

[26] JY He, PW Chan, QS Li, and CW Lee. 2022. Characterizing coastal wind energy resources based on sodar and microwave radiometer observations. *Renewable and Sustainable Energy Reviews* 163 (2022), 112498.

[27] Jürgen Hess, Maximilian Beinhofer, and Wolfram Burgard. 2014. A probabilistic approach to high-confidence cleaning guarantees for low-cost cleaning robots. In *2014 IEEE International Conference on Robotics and Automation (ICRA)*. 5600–5605. https://doi.org/10.1109/ICRA.2014.6907682

[28] InspectAPedia. 2022. Air Flow Rate in HVAC Systems CFM & fpm air flow speed data for building air ducts, air handlers, air conditioners & heating furnaces. https://inspectapedia.com/aircond/Air_Flow_Rates.php.

[29] Accutron Instruments. 2022. FlowTrax - Ultrasonic airflow monitor. https://accutroninstruments.com/flowtrax-ultrasonic-airflow-monitor/.

[30] Stefan Isler, Reza Sabzevari, Jeffrey A. Delmerico, and Davide Scaramuzza. 2016. An information gain formulation for active volumetric 3D reconstruction. In *2016 IEEE International Conference on Robotics and Automation, ICRA 2016, Stockholm, Sweden, May 16-21, 2016*, Danica Kragic, Antonio Bicchi, and Alessandro De Luca (Eds.). IEEE, 3477–3484. https://doi.org/10.1109/ICRA.2016.7487527

[31] Margarita A Kallistratova, Rostislav D Kouznetsov, Valerii F Kramar, and Dmitrii D Kuznetsov. 2013. Profiles of wind speed variances within nocturnal low-level jets observed with a sodar. *Journal of Atmospheric and Oceanic Technology* 30, 9 (2013), 1970–1977.

[32] Ryota Kikuchi, Takashi Misaka, Shigeru Obayashi, and Hamaki Inokuchi. 2020. Real-time estimation of airflow vector based on lidar observations for preview control. *Atmospheric Measurement Techniques* 13, 12 (2020), 6543–6558.

[33] Manikanta Kotaru, Kiran Joshi, Dinesh Bharadia, and Sachin Katti. 2015. SpotFi: Decimeter Level Localization Using WiFi. *SIGCOMM*

Comput. Commun. Rev. 45, 4 (aug 2015), 269–282. https://doi.org/10.1145/2829988.2787487

[34] Keun-Sang Kwon, Jung-Im Park, Young Joon Park, Don-Myung Jung, Ki-Wahn Ryu, and Ju-Hyung Lee. 2020. Evidence of long-distance droplet transmission of SARS-CoV-2 by direct air flow in a restaurant in Korea. *Journal of Korean medical science* 35, 46 (2020).

[35] S Lakshmi Priyadarsini and M Suresh. 2020. Factors influencing the epidemiological characteristics of pandemic COVID 19: A TISM approach. *International Journal of Healthcare Management* 13, 2 (2020), 89–98.

[36] Steven Lang and Eamon McKeogh. 2011. LIDAR and SODAR measurements of wind speed and direction in upland terrain for wind energy purposes. *Remote Sensing* 3, 9 (2011), 1871–1901.

[37] Yanchen Liu, Minghui Zhao, Stephen Xia, Eugene Wu, and Xiaofan Jiang. 2022. A sensorless drone-based system for mapping indoor 3D airflow gradients: demo abstract. In *Proceedings of the 20th Annual International Conference on Mobile Systems, Applications and Services*. 634–635.

[38] Giuseppe Loianno, Chris Brunner, Gary McGrath, and Vijay Kumar. 2016. Estimation, control, and planning for aggressive flight with a small quadrotor with a single camera and IMU. *IEEE Robotics and Automation Letters* 2, 2 (2016), 404–411.

[39] Jianyun Lu, Jieni Gu, Kuibiao Li, Conghui Xu, Wenzhe Su, Zhisheng Lai, Deqian Zhou, Chao Yu, Bin Xu, and Zhicong Yang. 2020. COVID-19 outbreak associated with air conditioning in restaurant, Guangzhou, China, 2020. *Emerging infectious diseases* 26, 7 (2020), 1628.

[40] Dani Martínez, Eduard Clotet, Marcel Tresanchez, Javier Moreno, Juan Manuel Jiménez-Soto, Rudys Magrans, Santiago Marco, and Jordi Palacín. 2015. First characterization results obtained in a wind tunnel designed for indoor gas source detection. In *2015 International Conference on Advanced Robotics (ICAR)*. IEEE, 629–634.

[41] D. McLean. 2012. *Understanding Aerodynamics: Arguing from the Real Physics*. Wiley. https://books.google.com/books?id=UE3sxu28R0wC

[42] Jennifer McWilliams. 2002. Review of air flow measurement techniques. (2002).

[43] L.M. Milne-Thomson. 1973. *Theoretical Aerodynamics*. Dover Publications. https://books.google.com/books?id=Az8C1Qp55GcC

[44] ModalAI. 2021. Seeker - Autonomous SLAM Indoor Drone. https://www.modalai.com/pages/seeker-autonomous-indoor-drone.

[45] Kartik Mohta, Michael Watterson, Yash Mulgaonkar, Sikang Liu, Chao Qu, Anurag Makineni, Kelsey Saulnier, Ke Sun, Alex Zhu, Jeffrey Delmerico, et al. 2018. Fast, autonomous flight in GPS-denied and cluttered environments. *Journal of Field Robotics* 35, 1 (2018), 101–120.

[46] Patrick P Neumann and Matthias Bartholmai. 2015. Real-time wind estimation on a micro unmanned aerial vehicle using its inertial measurement unit. *Sensors and Actuators A: Physical* 235 (2015), 300–310.

[47] Simon Ottenhaus, Lukas Kaul, Nikolaus Vahrenkamp, and Tamim Asfour. 2018. Active Tactile Exploration Based on Cost-Aware Information Gain Maximization. *Int. J. Humanoid Robotics* 15, 1 (2018), 1850015:1–1850015:21. https://doi.org/10.1142/S0219843618500159

[48] Nikolaos Passalis and Anastasios Tefas. 2020. Continuous drone control using deep reinforcement learning for frontal view person shooting. *Neural Computing and Applications* 32, 9 (2020), 4227–4238.

[49] IA Perez, MA Garcıa, ML Sanchez, and B De Torre. 2004. Analysis of height variations of sodar-derived wind speeds in Northern Spain. *Journal of Wind Engineering and Industrial Aerodynamics* 92, 10 (2004), 875–894.

[50] Gabriele Perozzi, Denis Efimov, Jean-Marc Biannic, and Laurent Planckaert. 2022. Using a quadrotor as wind sensor: Time-varying parameter estimation algorithms. *Internat. J. Control* 95, 1 (2022), 126–137.

[51] HALO Photonics. 2020. Streamline series. https://halo-photonics.com/lidar-systems/stream-line-series/.

[52] Julio A. Placed, Jared Strader, Henry Carrillo, Nikolay Atanasov, Vadim Indelman, Luca Carlone, and José A. Castellanos. 2022. A Survey on Active Simultaneous Localization and Mapping: State of the Art and New Frontiers. *CoRR* abs/2207.00254 (2022). https://doi.org/10.48550/arXiv.2207.00254 arXiv:2207.00254

[53] J.D. Posner, C.R. Buchanan, and D. Dunn-Rankin. 2003. Measurement and prediction of indoor air flow in a model room. *Energy and Buildings* 35, 5 (2003), 515–526. https://doi.org/10.1016/S0378-7788(02)00163-9

[54] D. Prakash and P. Ravikumar. 2015. Analysis of thermal comfort and indoor air flow characteristics for a residential building room under generalized window opening position at the adjacent walls. *International Journal of Sustainable Built Environment* 4, 1 (2015), 42–57. https://doi.org/10.1016/j.ijsbe.2015.02.003

[55] Anis Naema Atiyah Rafai, Noraziah Adzhar, and Nor Izzati Jaini. 2022. A Review on Path Planning and Obstacle Avoidance Algorithms for Autonomous Mobile Robots. *J. Robotics* 2022 (2022), 2538220:1–2538220:14. https://doi.org/10.1155/2022/2538220

[56] Ring. 2022. A new kind of camera. That flies. https://ring.com/always-home-cam-flying-camera.

[57] Indoor Robotics. 2022. Tando - Autonomous Indoor Drone. https://www.indoor-robotics.com/solution/.

[58] Andres Rodriguez, Evan Andersen, Justin Bradley, and Clark Taylor. 2007. Wind estimation using an optical flow sensor on a miniature air vehicle. In *AIAA Guidance, Navigation and Control Conference and Exhibit*. 6614.

[59] Kansuke Sasaki, Minoru Inoue, Tomoya Shimura, and Masato Iguchi. 2021. In Situ, Rotor-Based Drone measurement of wind vector and aerosol concentration in volcanic areas. *Atmosphere* 12, 3 (2021), 376.

[60] Nikolaus Peter Schmitt, Wolfgang Rehm, Thomas Pistner, Paul Zeller, Hermann Diehl, and Peter Navé. 2007. The AWIATOR airborne LIDAR turbulence sensor. *Aerospace Science and Technology* 11, 7-8 (2007), 546–552.

[61] Zhu Shi, Zechao Lu, and Qingyan Chen. 2019. Indoor airflow and contaminant transport in a room with coupled displacement ventilation and passive-chilled-beam systems. *Building and Environment* 161 (2019), 106244. https://doi.org/10.1016/j.buildenv.2019.106244

[62] Magdalena Simma, Håvard Mjøen, and Tobias Boström. 2020. Measuring wind speed using the internal stabilization system of a quadrotor drone. *Drones* 4, 2 (2020), 23.

[63] Reid Simmons, David Apfelbaum, Wolfram Burgard, Dieter Fox, Mark Moors, Sebastian Thrun, and Håkan Younes. 2000. Coordination for multi-robot exploration and mapping. In *Aaai/Iaai*. 852–858.

[64] Ewoud JJ Smeur, Qiping Chu, and Guido CHE De Croon. 2016. Adaptive incremental nonlinear dynamic inversion for attitude control of micro air vehicles. *Journal of Guidance, Control, and Dynamics* 39, 3 (2016), 450–461.

[65] Ewoud JJ Smeur, Guido CHE de Croon, and Qiping Chu. 2018. Cascaded incremental nonlinear dynamic inversion for MAV disturbance rejection. *Control Engineering Practice* 73 (2018), 79–90.

[66] Ezra Tal and Sertac Karaman. 2020. Accurate tracking of aggressive quadrotor trajectories using incremental nonlinear dynamic inversion and differential flatness. *IEEE Transactions on Control Systems Technology* 29, 3 (2020), 1203–1218.

[67] ArduPilot Development Team. 2023. Copter Position Control and Navigation. https://ardupilot.org/dev/docs/code-overview-copter-poscontrol-and-navigation.html.

[68] PX4 Development Team and Community. 2022. Controller Diagrams. https://docs.px4.io/main/en/flight_stack/controller_diagrams.html.

[69] William Thielicke, Waldemar Hübert, Ulrich Müller, Michael Eggert, and Paul Wilhelm. 2021. Towards accurate and practical drone-based wind measurements with an ultrasonic anemometer. *Atmospheric Measurement Techniques* 14, 2 (2021), 1303–1318.

[70] Sebastian Thrun, Wolfram Burgard, and Dieter Fox. 2005. *Probabilistic Robotics (Intelligent Robotics and Autonomous Agents).* The MIT Press.

[71] Sebastian Thrun, Scott Thayer, William Whittaker, Christopher Baker, Wolfram Burgard, David Ferguson, Dirk Hahnel, D Montemerlo, Aaron Morris, Zachary Omohundro, et al. 2004. Autonomous exploration and mapping of abandoned mines. *IEEE Robotics & Automation Magazine* 11, 4 (2004), 79–91.

[72] L Valldecabres, N Nygaard, L Vera-Tudela, L Von Bremen, and M Kühn. 2018. On the use of dual-doppler radar measurements for very short-term wind power forecasts, Remote Sens., 10, 1701.

[73] T. van Hooff and B. Blocken. 2010. Coupled urban wind flow and indoor natural ventilation modelling on a high-resolution grid: A case study for the Amsterdam ArenA stadium. *Environmental Modelling & Software* 25, 1 (2010), 51–65. https://doi.org/10.1016/j.envsoft.2009.07.008

[74] Madhu Babu Vankadari, Kaushik Das, Chinmay Shinde, and Swagat Kumar. 2018. A reinforcement learning approach for autonomous control and landing of a quadrotor. In *2018 International Conference on Unmanned Aircraft Systems (ICUAS).* IEEE, 676–683.

[75] Nikola Vasiljević, Michael Harris, Anders Tegtmeier Pedersen, Gunhild Rolighed Thorsen, Mark Pitter, Jane Harris, Kieran Bajpai, and Michael Courtney. 2020. Wind sensing with drone-mounted wind lidars: proof of concept. *Atmospheric Measurement Techniques* 13, 2 (2020), 521–536.

[76] Wei Wang and A. J. Roberts. 2015. Diffusion Approximation for Self-Similarity of Stochastic Advection in Burgers' Equation. *Communications in Mathematical Physics* 333, 3 (01 Feb 2015), 1287–1316. https://doi.org/10.1007/s00220-014-2117-7

[77] Curtis R Wood, L Pauscher, HC Ward, S Kotthaus, JF Barlow, M Gouvea, SE Lane, and CSB Grimmond. 2013. Wind observations above an urban river using a new lidar technique, scintillometry and anemometry. *Science of the total environment* 442 (2013), 527–533.

[78] Weirong Zhang, Yanan Zhao, Peng Xue, and Kunio Mizutani. 2022. Review and development of the contribution ratio of indoor climate (CRI). *Energy and Built Environment* 3, 4 (2022), 412–423.

[79] Minghui Zhao, Tyler Chang, Aditya Arun, Roshan Ayyalasomayajula, Chi Zhang, and Dinesh Bharadia. 2021. ULoc: Low-Power, Scalable and Cm-Accurate UWB-Tag Localization and Tracking for Indoor Applications. 5, 3, Article 140 (sep 2021), 31 pages. https://doi.org/10.1145/3478124

[80] Dapeng Zhou, Zeyu Jin, and Guoqiang Wu. 2023. Improved Adaptive NDI Flight Control Law Design Based on Real-Time Aerodynamic Identification in Frequency Domain. *Applied Sciences* 13, 12 (2023), 6951.