

# Experience-Based Models of Surface Proximal Aerial Robot Flight Performance in Wind

John W. Yao, Vishnu R. Desaraju, and Nathan Michael

Carnegie Mellon University,  
Pittsburgh, PA 15213, USA  
{johnyao, rajeswar, nmichael}@cmu.edu

**Abstract.** This work presents an experiment-driven aerodynamic disturbance modeling technique that leverages experiences from past flights to construct a predictive model of the exogenous forces acting on an aerial robot. Specifically, we consider operation in turbulent air stemming from the interaction between wind and nearby surfaces. To construct the disturbance model, we employ Locally Weighted Projection Regression and relate the aerial robot’s state and an experimentally learned freestream wind model to the disturbance forces estimated during flight through the environment. The approach is experimentally validated through a set of flight tests in an indoor environment with artificially generated turbulent airflow that illustrate the computation of this disturbance model, its generalizability across flow conditions, and its utility for disturbance-aware motion planning.

## 1 Introduction and Related Work

In this work, we construct models of aerial robot flight performance near surface obstacles under the influence of wind. Smooth and accurate flight is difficult to achieve in these conditions due to the complex, nonlinear effects of aerodynamic forces and surface-induced turbulence on vehicle motion. Consequently, applications such as aerial robot infrastructure inspection that require precise positioning of sensors to collect data are extremely sensitive to local aerodynamic forces. To mitigate performance degradation due to harmful aerodynamic effects and take advantage of beneficial ones, planning strategies rely on access to high fidelity models of aerodynamic disturbances throughout the flight environment. Therefore, building accurate models of flight performance with respect to varying wind and local surface geometry is an important capability to improve the quality of aerial robot infrastructure inspection.

Several approaches address the problem of planning and control with aerial robot plant models that vary with time and environmental conditions. One class of methods seeks to develop precise empirical models of aerodynamic phenomena [1–4] that can be accurately inverted to compute commands that achieve desired motions. Unfortunately, the weak observability [5] of many such models with available sensors in non-laboratory flight conditions renders them incapable of capturing changes in vehicle dynamics.

Adaptive control strategies reactively update the controller’s internal plant model to match the vehicle’s actual motion and are widely used in aerial robots to mitigate

unmodeled dynamics [6]. However, these methods are limited by the observation delay of inertial sensors and pose estimation [7]. Although sensors such as strain gages [8], accelerometer arrays [9], and pitot tubes [10] can reduce this delay and facilitate faster cancellation of disturbances, adaptive control strategies cannot pre-emptively mitigate them.

Anticipating aerodynamic disturbances in planning and control requires generative models of the impact of environmental phenomena on vehicle motion [11], such as wind velocity on gliders [12] and ocean currents on underwater robots [13]. In many cases, the disturbance field is strongly correlated with the geometry of the vehicle’s immediate surroundings [14]. This fact is exploited in works that infer aerodynamic disturbance forces from local geometry [15].

In this work, we propose and experimentally validate an approach for modeling aerodynamic disturbance forces on an aerial robot flying near surface obstacles in wind. The proposed method leverages an aerial robot’s accumulated flight experiences to predict disturbance forces that it will experience at a given position, velocity, and for a given set of far field wind conditions. We summarize the technical approach in Sect. 2. Section 3 presents the results of experiments that validate the functionality and accuracy of the proposed horizontal aerodynamic disturbance force model under a variety of trajectories and artificially generated wind conditions in an indoor laboratory setting. Additionally, we demonstrate the utility of this model for aerial robot trajectory planning near surface obstacles in strong winds. Section 4 closes with a discussion and summary of the implications of the experimental results.

## 2 Technical Approach

### 2.1 Aerodynamic Disturbance Estimation

In this work, we consider the translational dynamics of a quadrotor aerial robot with respect to a world reference frame,  $\{W\}$ , and a body reference frame,  $\{B\}$ , that is coincident with the center of mass such that the positive  $z$ -axis points in the nominal thrust direction parallel to the rotor axes.

$$\dot{\mathbf{p}} = \mathbf{v} \tag{1}$$

$$\dot{\mathbf{v}} = \frac{1}{m} (T\mathbf{R}\mathbf{e}_3 + \mathbf{f}_t + \mathbf{f}_a) - g\mathbf{e}_3 \tag{2}$$

The vectors  $\mathbf{p}, \mathbf{v} \in \{W\}$  are the position and velocity of the vehicle center of mass,  $\mathbf{R}$  is a rotation matrix that takes vectors from  $\{B\}$  to  $\{W\}$ ,  $\mathbf{e}_3$  is the third column of the  $3 \times 3$  identity matrix, and  $\mathbf{f}_t, \mathbf{f}_a \in \{W\}$  are the trims and aerodynamic disturbance forces acting on the vehicle. The constants  $m$  and  $g$  denote mass and gravitational acceleration, respectively, while the variable  $T$  represents the total thrust from all rotors along the body  $z$ -axis. We defer details of the lower-level attitude and motor dynamics to the relevant source literature [2].

We estimate the acceleration associated with the aerodynamic disturbance force,  $\mathbf{a} = \frac{\mathbf{f}_a}{m}$ , using a nonlinear observer that accounts for motor dynamics:

$$\dot{\hat{\mathbf{v}}} = \frac{c_T}{m} \hat{\boldsymbol{\omega}}^T \hat{\boldsymbol{\omega}} \mathbf{R} \mathbf{e}_3 - g \mathbf{e}_3 + \hat{\mathbf{a}} + \frac{\mathbf{f}_t}{m} + \mathbf{L} (\mathbf{v} - \hat{\mathbf{v}}) \quad (3)$$

$$\dot{\hat{\boldsymbol{\omega}}} = k_m (\boldsymbol{\omega}^d - \hat{\boldsymbol{\omega}}) \quad (4)$$

$$\dot{\hat{\mathbf{a}}} = \mathbf{\Gamma} (\mathbf{v} - \hat{\mathbf{v}}) \quad (5)$$

The vectors  $\hat{\mathbf{v}}$ ,  $\hat{\boldsymbol{\omega}}$ , and  $\hat{\mathbf{a}}$  are estimates of velocity, rotor speeds, and aerodynamic disturbance acceleration. Observations of velocity and attitude are provided by a separate state estimator, while the desired rotor speeds,  $\boldsymbol{\omega}^d$ , are provided by the controller. Trims, mass, thrust coefficient ( $c_T$ ), and motor constant ( $k_m$ ) are measured offline in separate experiments.  $\mathbf{L}$  and  $\mathbf{\Gamma}$  are diagonal matrices of observer gains. The aerodynamic disturbance acceleration observation,  $\hat{\mathbf{a}}$ , is obtained by low pass filtering  $\hat{\mathbf{a}}$  to neglect high frequency components that are beyond the bandwidth of the vehicle's actuation.

## 2.2 Freestream Wind Descriptor

We define the freestream wind descriptor,  $\mathbf{a}_\infty$ , as the aerodynamic disturbance that would be experienced by a vehicle at a location in a wind field in the absence of all surface obstacles. As this measure can account for the majority of bulk airflow effects, other features such as position and velocity can be used to model local aerodynamic disturbance variations from surface-induced effects. Additionally, this descriptor allows for comparison and interpolation between different bulk airflow conditions (see Sect. 3.2).

To evaluate the freestream wind descriptor, the aerial robot traverses a closed path trajectory offset from the obstacles at a distance such that the experienced aerodynamic disturbances along the path are primarily due to the bulk airflow and not surface-induced effects. The disturbance accelerations observed along the path are introduced into a Gaussian process regression [16] to create a mapping from position to disturbance acceleration

$$\mathbf{a}_\infty(\mathbf{p}) = \mathcal{GP}(\mathbf{p}) \quad (6)$$

To account for the dominant direction of wind flow, the Gaussian process distance measure is weighted according to the mean disturbance acceleration in each direction. Figure 1a illustrates the  $y$ -component of the freestream wind descriptor model computed from a circular trajectory in a unidirectional wind field, leading to comparable results computed over a trajectory that passes through interior locations (Fig. 1b). The similarity of these results indicates that the information obtained from the perimeter of the environment provides insight into the true aerodynamic disturbances in the interior, validating the proposed method of computing the freestream wind descriptor.

## 2.3 Aerodynamic Disturbance Prediction

Locally Weighted Projection Regression (LWPR) [17] is employed to generate an aerodynamic disturbance prediction model. LWPR is a data-efficient regression technique

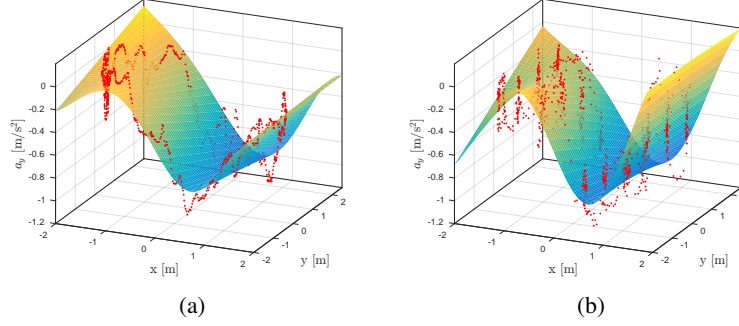


Fig. 1: The freestream wind descriptor computed via (a) perimeter data and (b) interior data exhibit similar predictions for the  $y$ -component of disturbance acceleration over the  $x$  and  $y$ -components of position.

that computes a low-dimensional model of the training data based on a combination of local linear models. Data efficiency is particularly important in an experience-based approach to enable continual learning and reuse of learned models.

The disturbance model inputs are the aerial robot’s position, velocity, and the freestream wind descriptor. Separate LWPR models (7) are learned to predict each component of the disturbance acceleration observed via (3)–(5)

$$\mu_{a,i} = l_i(\mathbf{p}, \mathbf{v}, \mathbf{a}_\infty(\mathbf{p})), \quad i = \{x, y, z\} \quad (7)$$

In addition to the mean prediction,  $\mu_{a,i}$ , LWPR also provides a standard deviation uncertainty bound,  $\sigma_{a,i}$ . For notational convenience, we stack the componentwise means and uncertainties into the vectors  $\boldsymbol{\mu}_a$  and  $\boldsymbol{\sigma}_a$ . While the formulation presented in this section is applicable to three dimensions, the experiments detailed in the sequel use two dimensional positions, velocities, and disturbance accelerations for ease of visualization.

### 3 Experiments

In this section, we detail two experimental studies that seek to validate the aerodynamic prediction model developed above and illustrate its utility for disturbance-aware motion planning.

#### 3.1 Experimental Setup and Data Collection

To achieve repeatable wind conditions for model validation, we conducted experiments in an indoor  $5 \times 5 \times 4 \text{ m}^3$  net-enclosed flight arena surrounded by movable high-power fans that generate turbulent airflow of up to 6.0 m/s. To disrupt the wind field and generate realistic aerodynamic interactions with surfaces, we introduce a transparent cube in the center of the flight environment. Flight tests are performed on a Lumenier Danaus

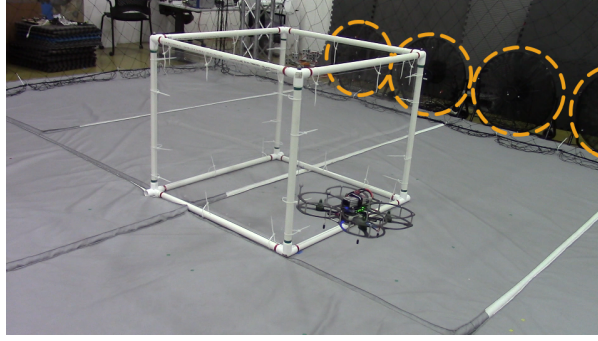


Fig. 2: The quadrotor aerial robot navigates around a transparent cube obstacle in a wind field generated by high-power fans (orange circles), each of which is 70 cm in diameter and capable of producing turbulent airflow conditions up to 6.0 m/s.

quadrotor equipped with a Pixhawk autopilot and Odroid U3 computer. The vehicle fits within a  $0.29 \times 0.32 \times 0.14 \text{ m}^3$  volume and has a hover mass of 0.734 kg. An unscented Kalman filter [18] fuses IMU data with motion capture pose observations to provide a smooth state estimate for control. A cascaded control scheme [19] consisting of a 100 Hz outer position controller and a 200 Hz inner attitude controller, both augmented with L1 adaptive control [20] for disturbance rejection, is implemented in C++ using ROS and run onboard the vehicle. Figure 2 illustrates the experimental platform and environment.

To obtain a diverse set of wind conditions for building the aerodynamic disturbance model, different airflow patterns are synthesized by aligning a row of high-powered fans to point east, south, and southeast while varying their strength. Freestream wind models are computed for each wind condition using data from circular flights along the perimeter of the flight arena. Next, the aerial robot is commanded to fly trajectories around the cube obstacle at a constant height of 60 cm above the ground. During postprocessing, the position, velocity, disturbance acceleration, and freestream wind descriptor (computed from position using (6)) from the vehicle's flight logs are used to train the aerodynamic disturbance model.

### 3.2 Model Prediction Performance

The utility of the aerodynamic disturbance model is contingent on its ability to generate reasonable predictions in wind conditions that differ from those encountered in previous flights, as is often the case when a vehicle returns to inspect a structure after a long time interval. We assume that the bulk airflow is slowly varying or constant over the spatiotemporal domain of a typical inspection flight mission. Therefore, the freestream wind descriptor provides a way to express the “distance” between two bulk flow conditions. This component of the proposed aerodynamic disturbance model enables generalization across different wind conditions.

To demonstrate this fact, we conducted twelve flight tests around the cube obstacle under westerly and northerly wind of varying strengths and one flight under northwest-

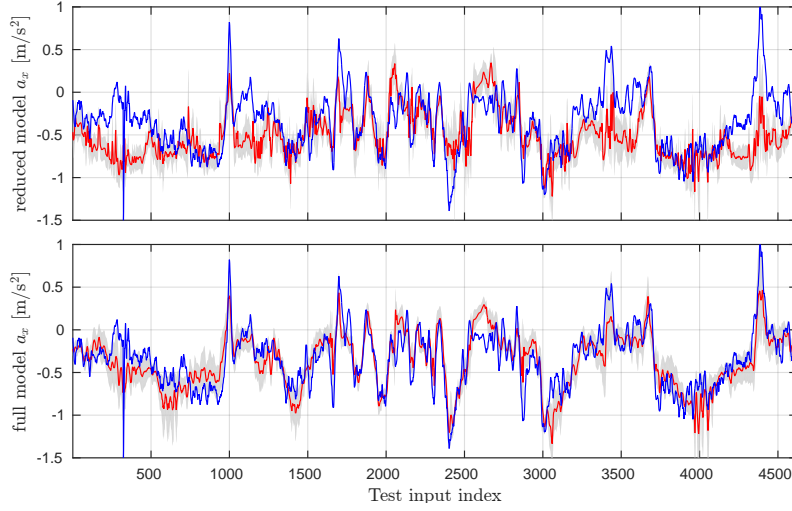


Fig. 3: The actual (blue), predicted (red), and  $1\sigma$  uncertainty envelope (gray) for the  $x$ -component of the aerodynamic disturbance acceleration over the test trajectory under northwesterly wind are depicted for the reduced (top) and full (bottom) models.

erly wind. Two LWPR aerodynamic disturbance models are trained on data from the first group of flights. The reduced model’s input feature vector consists of only position and velocity, while the full model’s input feature vector also includes the freestream wind descriptor. Both models are tested on flight data from the single flight under northwesterly wind. A comparison of both models’ predictions for  $a_x$  against the ground truth (Fig. 3) shows that the full model is more accurate than the reduced model. Figure 4 supports this conclusion and shows that the full model’s prediction error norm (magnitude of  $a_x$  and  $a_y$  prediction errors) over the test data is lower than that of the reduced model. The oracle model uses the same LWPR hyperparameters as the full model, but is trained on the *test* data, and is included in Fig. 4 for comparison as the benchmark for the best achievable prediction accuracy under the same LWPR hyperparameters.

The superior performance of the full model with respect to the reduced model can be attributed to the fact that the latter only encodes position and velocity data to account for variations in bulk flow, resulting in degradation of prediction accuracy as training data is drawn from an increasing variety of wind conditions. The results in Figs. 3 and 4 highlight the improved prediction of disturbance accelerations in a previously unseen northwesterly wind using only flight data in westerly and northerly winds and demonstrates the benefit of using the freestream wind descriptor as a means of generalizing across different bulk flow conditions.

### 3.3 Disturbance-Aware Motion Planning

To demonstrate the utility of the proposed aerodynamic disturbance prediction model, we consider a scenario derived from persistent inspection applications in which a quadrotor must navigate around a structure in the presence of an initially unknown wind field.

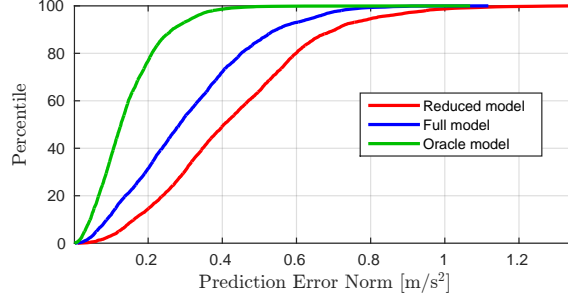


Fig. 4: The cumulative distributions of disturbance acceleration prediction error norms for the full, reduced, and oracle models over a test flight under northwesterly wind.

However, as the inspection task requires multiple traversals of the region, we can leverage the disturbance model learned from past experience to improve the traversal route.

An A\* search with an Euclidean distance heuristic computes optimal paths between user-selected start and goal positions. The resulting path is converted into a smooth trajectory for the quadrotor to follow by fitting a minimum-acceleration polynomial spline [21]. Although we restrict plans to the 2D horizontal plane for simplicity, the proposed approach generalizes to 3D. The plane is discretized with each grid cell connected to each of its eight neighbors, while each edge can be traversed at three different speeds. The cost,  $J$ , at each neighbor is defined as

$$J(\mathbf{p}, \mathbf{v}, \mathbf{a}_\infty) = \Delta\mathbf{p} + h(\mathbf{p}, \mathbf{v}, \mathbf{a}_\infty) \quad (8)$$

$$h(\mathbf{p}, \mathbf{v}, \mathbf{a}_\infty) = \|\boldsymbol{\mu}_a\| + \|\boldsymbol{\sigma}_a\| \quad (9)$$

where  $\Delta\mathbf{p}$  is the distance between the neighboring cell and the current one. We define field cost ( $h$ ) as the sum of the vector norms of the predicted acceleration mean and uncertainty, penalizing strong disturbances as well as volatile regions. However, as LWPR reports high uncertainty in regions with sparse flight data coverage, we only apply the field cost when the prediction model uncertainty is within empirically determined bounds on model variability.

At the beginning of each planning trial, the quadrotor computes a path to a location on the other side of a cube obstacle with a strong wind field approximately orthogonal to the desired direction of travel, as depicted in Fig. 5. A series of five successive trials demonstrates the evolution of the computed trajectory as the accumulation of flight experience improves the accuracy of the field cost.

In the first planning trial, the vehicle starts with a prior that encodes a maximally uncertain aerodynamic disturbance model and plans a trajectory that ignores the field cost. After each trial, the model is updated according to the recorded observations and the task is repeated using the growing set of accumulated experiences. Figure 6 shows the computed trajectories for all trials and the updated field costs for all trials excluding the first trial.

In the second, third, and fourth trials (Figs. 6a-6c), the trajectory alternates between the windward and leeward sides of the obstacle to reduce uncertainty in the aerody-

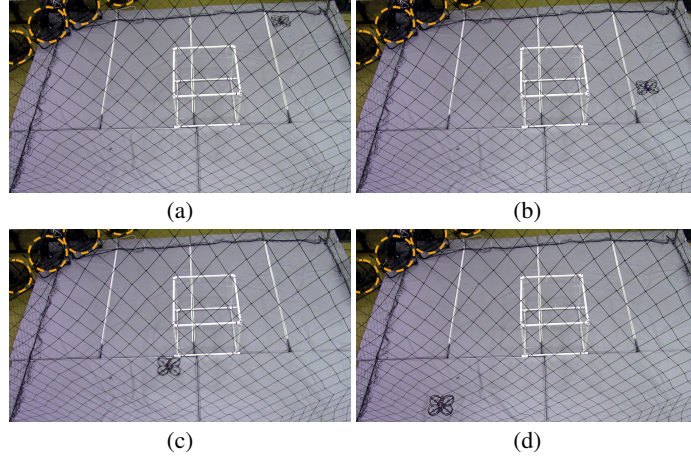


Fig. 5: Snapshots of the quadrotor executing the disturbance-aware path computed in the final planning trial. Fans (orange circles) generate a northwesterly bulk flow.

dynamic model. In the final trial (Fig. 6d), the planner converges to a trajectory on the leeward side of the obstacle, shielded from direct wind from the fans while avoiding the turbulent regions behind the obstacle (Fig. 6c).

Table 1 reports an overall trend of decreasing aerodynamic disturbance acceleration magnitudes and variability as more accurate field costs are used for planning. As a consequence of the planner explicitly optimizing for low disturbance trajectories, both mean and maximum tracking error are reduced over the sequence of runs. By utilizing the proposed experience-based aerodynamic disturbance model, the quadrotor is able to compute informed paths through the environment that minimize adverse aerodynamic effects of the wind field on flight performance.

## 4 Conclusion

In this work, we investigate the problem of modeling aerodynamic disturbances on aerial robots operating in windy environments near obstacles. We present a modeling

Table 1: Horizontal disturbance acceleration norm and 3D tracking error norm over successive A\* planner trajectories.

Trial	Horizontal Disturbance Norm [m/s <sup>2</sup> ]		3D Tracking error norm [m]	
	Mean	Standard Deviation	Mean	Max
1	0.305	0.342	0.065	0.460
2	0.276	0.276	0.069	0.427
3	0.286	0.336	0.065	0.418
4	0.221	0.233	0.061	0.508
5	0.209	0.218	0.051	0.352

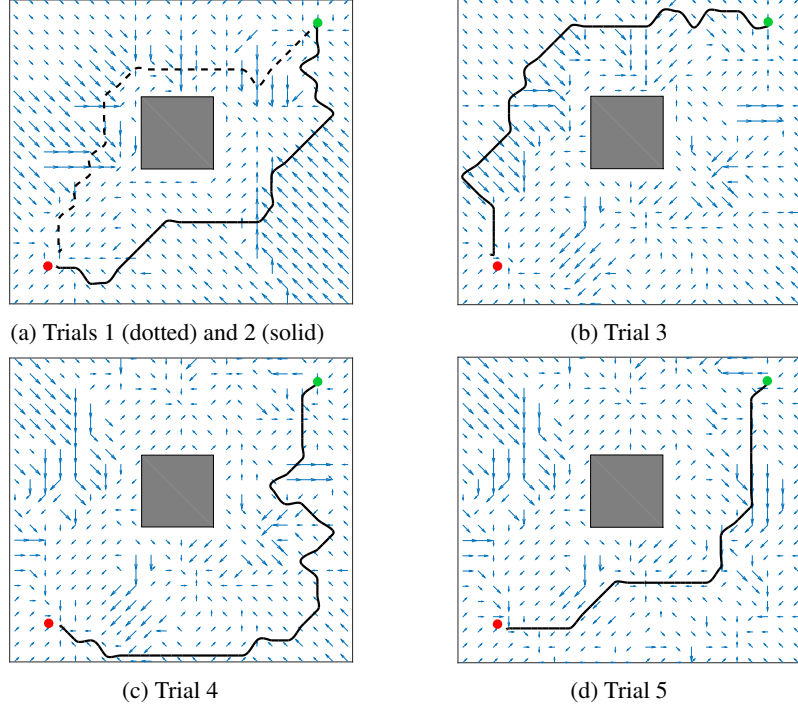


Fig. 6: The 2D paths (black) computed by the A\* planner on successive trials are shown relative to the central obstacle (gray) for the wind conditions depicted in Fig. 5. Start (green) and goal (red) locations correspond to the aerial robot's location in Fig. 5a and Fig. 5d, respectively. The vector arrows (blue) indicate the speed and direction of the motion that incurs the lowest field cost (9) at each grid cell in the environment.

technique that leverages past flight experiences to learn a mapping from position, velocity, and freestream wind conditions to disturbance acceleration. The experimental validation of these techniques suggest that the proposed methodology yields superior flight performance for surface-proximal operations with variable wind flow conditions. We demonstrate the viability of this approach in a laboratory setting with repeatable turbulent wind conditions and show that the approach leads to improvements in trajectory planning with the accumulation of flight experience over multiple missions, enabling informed motion planning to mitigate the effects of aerodynamic disturbances.

## References

1. Hoffmann, G.M., Huang, H., Waslander, S.L., Tomlin, C.J.: Quadrotor helicopter flight dynamics and control: Theory and experiment. In: Proc. of the AIAA Guidance, Navigation, and Control Conf. Volume 2., Hilton Head, USA (August 2007)
2. Bangura, M., Mahony, R.: Nonlinear dynamic modeling for high performance control of a quadrotor. In: Proc. of the Aus. Conf. on Robot. and Autom., Wellington, New Zealand (December 2012) 1–10

3. Omari, S., Hua, M.D., Ducard, G., Hamel, T.: Nonlinear control of vtol uavs incorporating flapping dynamics. In: Proc. of the IEEE/RSJ Intl. Conf. on Intell. Robots and Syst., Tokyo, Japan (November 2013) 2419–2425
4. Leishman, R., Macdonald, J., Beard, R., McLain, T.: Quadrotors and accelerometers: State estimation with an improved dynamic model. *IEEE Control Syst. Mag.* **34**(1) (February 2014) 28–41
5. Abeywardena, D., Wang, Z., Dissanayake, G., Waslander, S., Kodagoda, S.: Model-aided state estimation for quadrotor micro air vehicles amidst wind disturbances. In: Proc. of the IEEE/RSJ Intl. Conf. on Intell. Robots and Syst., Chicago, USA (September 2014) 4813–4818
6. Grande, R.C., Chowdhary, G., How, J.P.: Experimental validation of bayesian nonparametric adaptive control using gaussian processes. *J. Aero. Inf. Syst.* **11**(9) (July 2014) 565–578
7. Mohamed, A., Abdulrahim, M., Watkins, S., Clothier, R.: Development and flight testing of a turbulence mitigation system for micro air vehicles. *J. Field Robot.* **33**(5) (August 2016) 639–660
8. Ranganathan, B.N., Penskiy, I., Dean, W., Bergbreiter, S., Humbert, J.S.: Bio-inspired wind frame state sensing and estimation for mav applications. In: Proc. of the IEEE/RSJ Intl. Conf. on Intell. Robots and Syst., Hamburg, Germany (September 2015) 2729–2735
9. Gremillion, G.M., Castano, L.M., Humbert, J.S.: Disturbance rejection with distributed acceleration and strain sensing. In: Proc. of the AIAA Guidance, Navigation, and Control Conf. Volume 2., Kissimmee, USA (January 2015) 1623–1639
10. Yeo, D.W., Sydney, N., Paley, D.A.: Onboard flow sensing for rotary-wing uav pitch control in wind. In: Proc. of the AIAA Guidance, Navigation, and Control Conf. Volume 4., National Harbor, USA (January 2016) 2445–2455
11. Desaraju, V., Michael, N.: Hierarchical adaptive planning in environments with uncertain, spatially-varying disturbance forces. In: Proc. of the IEEE Intl. Conf. on Robot. and Autom., Hong Kong, China (May 2014) 5171–5176
12. Lawrance, N., Sukkarieh, S.: Path planning for autonomous soaring flight in dynamic wind fields. In: Proc. of the IEEE Intl. Conf. on Robot. and Autom., Shanghai, China (May 2011) 2499–2505
13. Hollinger, G.A., Pereira, A.A., Binney, J., Somers, T., Sukhatme, G.S.: Learning uncertainty in ocean current predictions for safe and reliable navigation of underwater vehicles. *J. Field Robot.* **33**(1) (January 2016) 47–66
14. Galway, D., Etele, J., Fusina, G.: Modeling of urban wind field effects on unmanned rotorcraft flight. *J. Aircraft* **48**(5) (September 2011) 1613–1620
15. Bartholomew, J., Calway, A., Mayol-Cuevas, W.: Learning to predict obstacle aerodynamics from depth images for micro air vehicles. In: Proc. of the IEEE Intl. Conf. on Robot. and Autom., Hong Kong, China (May 2014) 4967–4973
16. Rasmussen, C.E., Williams, C.K.I.: *Gaussian Processes for Machine Learning*. MIT Press (2006)
17. Vijayakumar, S., D’Souza, A., Schaal, S.: Incremental online learning in high dimensions. *Neural Computation* **17**(12) (December 2005) 2602–2634
18. Julier, S.J., Uhlmann, J.K.: A new extension of the kalman filter to nonlinear systems. In: Proc. SPIE. Volume 3068., Orlando, USA (July 1997) 182–193
19. Mahony, R., Kumar, V., Corke, P.: Multirotor aerial vehicles: Modeling, estimation, and control of quadrotor. *IEEE Robot. Autom. Mag.* **19**(3) (September 2012) 20–32
20. Hovakimyan, N., Cao, C.: *L1 adaptive control theory: guaranteed robustness with fast adaptation*. Volume 21. SIAM (2010)
21. Richter, C., Bry, A., Roy, N.: Polynomial trajectory planning for aggressive quadrotor flight in dense indoor environments. In: Proc. of the Intl. Sym. of Robot. Research, Singapore (December 2013)