

Document Version

Final published version

Citation (APA)

Das, A., Yadav, R. D., Sun, S., Sun, M., Kaski, S., & Pan, W. (2025). DroneDiffusion: Robust Quadrotor Dynamics Learning with Diffusion Models. In *Proceedings of the IEEE International Conference on Robotics and Automation, ICRA 2025* (pp. 1604-1610). (Proceedings - IEEE International Conference on Robotics and Automation). IEEE. <https://doi.org/10.1109/ICRA55743.2025.11127523>

Important note

To cite this publication, please use the final published version (if applicable). Please check the document version above.

Copyright

In case the licence states "Dutch Copyright Act (Article 25fa)", this publication was made available Green Open Access via the TU Delft Institutional Repository pursuant to Dutch Copyright Act (Article 25fa, the Taverne amendment). This provision does not affect copyright ownership. Unless copyright is transferred by contract or statute, it remains with the copyright holder.

Sharing and reuse

Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

Takedown policy

Please contact us and provide details if you believe this document breaches copyrights. We will remove access to the work immediately and investigate your claim.

**Green Open Access added to [TU Delft Institutional Repository](#)
as part of the Taverne amendment.**

More information about this copyright law amendment
can be found at <https://www.openaccess.nl>.

Otherwise as indicated in the copyright section:
the publisher is the copyright holder of this work and the
author uses the Dutch legislation to make this work public.

DroneDiffusion: Robust Quadrotor Dynamics Learning with Diffusion Models

Avirup Das^{*1}, Rishabh Dev Yadav^{*1}, Sihao Sun², Mingfei Sun¹, Samuel Kaski^{1,3}, Wei Pan¹

Abstract—An inherent fragility of quadrotor systems stems from model inaccuracies and external disturbances. These factors hinder performance and compromise the stability of the system, making precise control challenging. Existing model-based approaches either make deterministic assumptions, utilize Gaussian-based representations of uncertainty, or rely on nominal models, all of which often fall short in capturing the complex, *multimodal* nature of real-world dynamics. This work introduces *DroneDiffusion*, a novel framework that leverages conditional diffusion models to learn quadrotor dynamics, formulated as a sequence generation task. *DroneDiffusion* achieves superior generalization to unseen, complex scenarios by capturing the temporal nature of uncertainties and mitigating error propagation. We integrate the learned dynamics with an adaptive controller for trajectory tracking with stability guarantees. Extensive experiments in both simulation and real-world flights demonstrate the robustness of the framework across a range of scenarios, including unfamiliar flight paths and varying payloads, velocities, and wind disturbances.

Project page: <https://sites.google.com/view/dronediffusion>.

I. INTRODUCTION

Robust and reliable control of quadrotors is crucial for their expanding role in real-world applications, ranging from autonomous inspection to agile maneuvering [1], [2]. Achieving this requires designing accurate dynamic models that facilitate precise trajectory tracking and maneuver execution [3]. However, the development of such models is hindered by the complex interplay of aerodynamic effects, unmodeled dynamics, and external disturbances. Despite the use of uncertainty bounds and adaptation laws, model-based control approaches [4], [5] remain difficult to implement due to the poorly characterized and stochastic uncertainties inherent in real-world quadrotor environments.

Recent efforts have leveraged Gaussian Processes (GPs) to model aerodynamic effects [3], [6], unknown disturbances [7], and unmodeled dynamics [8]. While these methods provide a convenient data-driven strategy, the non-parametric nature of the GP-based model leads to high computational complexity, which increases with the dataset

R. Y and W. Pan acknowledge travel support from the European Union’s Horizon 2020 research and innovation programme under Grant Agreement No 951847.

¹ Department of Computer Science, The University of Manchester, United Kingdom ({avirup.das, rishabh.yadav}@postgrad.manchester.ac.uk, {mingfei.sun,samuel.kaski,wei.pan}@manchester.ac.uk).

² Department of Cognitive Robotics, Delft University of Technology, Netherlands (s.sun-2@tudelft.nl).

³ Department of Computer Science, Aalto University, Finland.

* Equal contribution.

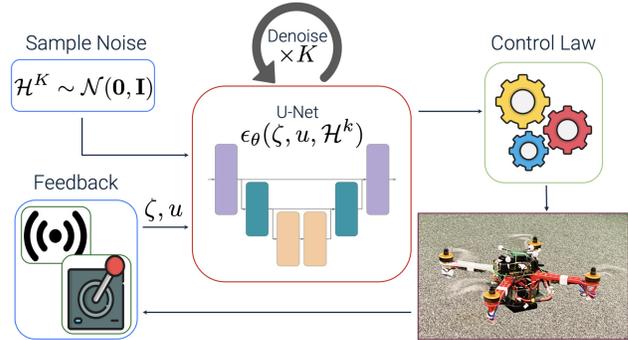


Fig. 1. DroneDiffusion leverages a diffusion model trained on sequences of sensory (state) data (ζ), control inputs (u), and corresponding unified residual dynamics terms \mathcal{H} to learn the conditional distribution $p(\mathcal{H}|\zeta, u)$. During flight, the framework uses a U-Net based noise predictor ϵ_θ to denoise \mathcal{H}^K (sampled from standard Gaussian $\mathcal{N}(\mathbf{0}, \mathbf{I})$) in K steps using (ζ, u) to generate a prediction of $\hat{\mathcal{H}}$. $\hat{\mathcal{H}}$ is then used by a nonlinear controller in a receding-horizon fashion for closed-loop feedback control of the quadrotor.

size, requiring the careful selection of representative training points [7]. As an alternative, Deep Neural Networks (DNNs) have been explored for system identification [9]–[11], with improvements from Spectrally Normalized DNNs [12]. However, despite their advantages, DNNs establish a deterministic mapping from the sensory inputs to the uncertain dynamics, limiting their adaptability to real-world flight conditions. Sequence models like Recurrent Neural Networks (RNNs) have also been effective in modeling dynamics by capturing long-range dependencies through multistep predictions [13], [14]. However, they require proper state initialization to avoid vanishing gradients. Transformers (e.g. GPT [15]) offer improvements over RNNs in capturing long-range dependencies and overall robustness but remain autoregressive in nature, which poses challenges in dynamics modeling. Autoregressive predictions can lead to cascading errors, where inaccuracies in initial predictions propagate through the sequence, causing progressively larger deviations over time [16], [17].

In real-world applications, the uncertainty in quadrotor dynamics is often governed by latent and difficult-to-measure factors such as airflow and downwash [18]–[20]. Developing precise nominal models to capture system dynamics presents significant challenges, necessitating a comprehensive knowledge of system parameters [21], [22] and limiting the ability of these methods to adapt to varying environmental conditions. Moreover, existing data-driven approaches [3], [6]–[8] often rely on deterministic models or assume predefined Gaussian structures for uncertainty, which limits their

capacity to capture the full complexity of the dynamics. Additionally, these methods frequently rely on partial uncertainty estimates derived from prior system dynamics information [12], [13], further constraining their adaptability to real-world scenarios. A toy experiment with a one-dimensional system and heavy-tailed disturbance (Figure 2) illustrates how standard methods struggle to capture the full distribution of dynamics, highlighting the need for robust models that do not rely on strong assumptions about the underlying dynamics and structure of disturbances.

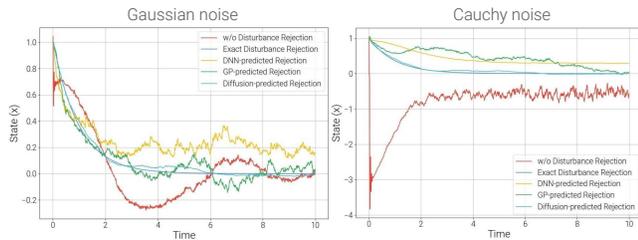


Fig. 2. Convergence comparison of the system $\dot{x}(t) = ax(t) + u + d$ with 8s of training data at 100Hz. Control inputs are generated by: $u = -Kx - \dot{d}$, where $a = 1$ and $K = 2$. *Left*: d is Gaussian. *Right*: d is a Cauchy disturbance. GP performs well in the presence of Gaussian disturbance but struggles with heavy-tailed disturbances due to kernel limitations. DNN converges with an offset for heavy-tailed disturbances but struggles to converge effectively when faced with Gaussian disturbance.

To address these limitations, we propose using diffusion models [23]–[25] for quadrotor dynamics modeling, which have recently demonstrated exceptional performance in fields such as image generation [26] and offline reinforcement learning [17], [27]. In this work, we apply diffusion models for the first time to model the complex, multimodal distributions present in quadrotor flight dynamics. Leveraging their ability to capture intricate data distributions and generate high-fidelity samples, we introduce **DroneDiffusion**, a framework that models quadrotor dynamics as a *conditional probabilistic distribution* and integrates it into a *nonlinear control framework*. The key contributions of this paper can be summarized as follows:

- We formulate dynamic model learning as a conditional sequence generation task using diffusion models, accurately characterizing the multimodal and stochastic nature of quadrotor dynamics in the real world.
- Our approach integrates the learned dynamic model with adaptive control for trajectory tracking, ensuring stability and adaptability to varying masses and velocities without extensive retraining.

Comprehensive real-time experiments validate the robustness of our framework in the face of model uncertainties, demonstrating its ability to accurately track complex trajectories, adapt to dynamic variations, and maintain resilience against wind disturbances, achieving significant advancements over existing solutions.

II. SYSTEM IDENTIFICATION WITH DIFFUSION MODELS

A. Preliminaries on diffusion models

Denosing Diffusion Probabilistic Models (DDPMs) [24] operate through forward and reverse processes. The forward

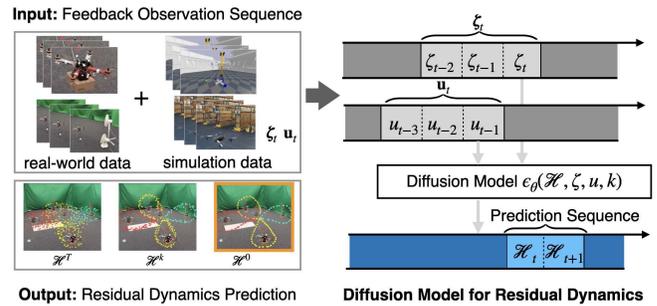


Fig. 3. Diffusion Model for Dynamics Learning. A general formulation (see Section II-B). At time step t , the residual dynamics takes the latest state and control input observation data ζ_t and u_t as input and outputs a sequence of residual dynamics prediction $[\mathcal{H}_t, \mathcal{H}_{t+1}, \dots]$.

process, $q(\mathbf{z}^k | \mathbf{z}^{k-1}) := \mathcal{N}(\mathbf{z}^k | \sqrt{1 - \beta^k} \mathbf{z}^{k-1}, \beta^k \mathbf{I})$, progressively transforms the data distribution $q(\mathbf{z}^0)$ into a standard Gaussian $\mathcal{N}(\mathbf{0}, \mathbf{I})$ over K steps, using a predefined variance schedule $\{\beta^k\}_{k=1}^K$. The reverse process, $p_\theta(\mathbf{z}^{k-1} | \mathbf{z}^k) := \mathcal{N}(\mathbf{z}^{k-1} | \mu_\theta(\mathbf{z}^k, k), \Sigma^k)$, parameterized by θ , reconstructs the data by estimating the mean μ_θ . Training optimizes a variational lower bound on the likelihood $\log p_\theta$ via the following objective:

$$\mathcal{J}_{\text{denoise}}(\theta) := \mathbb{E}_{k \sim \mathcal{U}(1, K), \mathbf{z}^0 \sim q, \epsilon \sim \mathcal{N}(\mathbf{0}, \mathbf{I})} \left[\|\epsilon - \epsilon_\theta(\mathbf{z}^k, k)\|^2 \right], \quad (1)$$

where $\epsilon_\theta(\mathbf{z}^k, k)$, a neural network parameterized by θ , predicts the noise ϵ injected during the forward process, enabling the recovery of \mathbf{z}^{k-1} . The reverse process can be conditioned on additional variables c by modifying μ_θ to incorporate c , resulting in a conditional reverse process: $p_\theta(\mathbf{z}^{k-1} | \mathbf{z}^k, c) := \mathcal{N}(\mathbf{z}^{k-1} | \mu_\theta(\mathbf{z}^k, c, k), \Sigma^k)$.

While DDPMs are commonly applied to image generation, we adapt them to model quadrotor dynamics, inspired by Diffusion Policy [28]. Our formulation introduces two key modifications: (1) redefining the output to represent quadrotor residual dynamics and (2) conditioning the denoising process on state observations ζ_t and past control inputs u_t . These changes are discussed below, with an overview provided in Figure 3.

B. Dynamic Model Learning as Sequence Generation

The position dynamics of a quadrotor [12] are governed by the states: position $p \triangleq [p_x, p_y, p_z]^\top \in \mathbb{R}^3$, velocity $v \in \mathbb{R}^3$, and attitude rotation matrix $R \in \text{SO}(3)$ as follows:

$$\dot{p} = v, \quad m\dot{v} = mg_v + Rf_u + f_a, \quad (2)$$

where m is the mass, $g_v = [0, 0, -g]^\top$ is the gravity vector, $f_u = [0, 0, T]^\top$ represents thrust, and f_a encapsulates unknown aerodynamic forces. To model the uncertain dynamics without assuming prior structural knowledge of f_a , we introduce a nominal mass $\bar{m} \in \mathbb{R}^+$, reformulating (2) as:

$$\bar{m}\ddot{p} + \mathcal{H}(p, \dot{p}, \ddot{p}) = u, \quad (3)$$

where $\mathcal{H}(p, \dot{p}, \ddot{p}) \triangleq (m - \bar{m})\ddot{p} - mg_v - f_a$, and $u = Rf_u$. The term \mathcal{H} captures state-dependent uncertainties arising due to aerodynamic effects f_a (e.g., drag, ground effect,

wind) and imperfect information of mass m (e.g., payloads), into a single variable. This reformulation facilitates a data-driven approach to infer the distributional properties of \mathcal{H} , which is inherently nonlinear and complex. Unlike single-step dynamic models, which often lack the temporal context required to infer the underlying causal structure of the dynamics and are prone to compounding errors, our approach frames the estimation of \mathcal{H} as a sequence generation task. This allows us to mitigate error propagation and make temporally consistent predictions.

For a time-horizon H , we define a sequence \mathcal{S} comprising uncertainties \mathcal{H} , control inputs u , and sensory measurements $\zeta = \{p, \dot{p}, \ddot{p}\}$:

$$\mathcal{S} = \begin{bmatrix} \zeta_1 & \zeta_2 & \dots & \zeta_H \\ u_1 & u_2 & \dots & u_H \\ \mathcal{H}_1 & \mathcal{H}_2 & \dots & \mathcal{H}_H \end{bmatrix}. \quad (4)$$

This induces a tight coupling between \mathcal{H} and (ζ, u) that makes the estimation of \mathcal{H} equivalent to sampling from the *conditional distribution* $p(\mathcal{H}|\zeta, u)$. By predicting \mathcal{H} over a horizon, we achieve temporally consistent predictions, ensuring a stable flight. During flight, the control input at the current timestep u_t is unavailable for estimating \mathcal{H} , so we use the control input from the previous timestep u_{t-1} , assuming one-step differences of control signals are small and bounded [12]. For simplicity, we retain the original notation in (4). At the first timestep, the previous control input is taken as zero.

C. Diffusion Model for Multimodal Dynamics

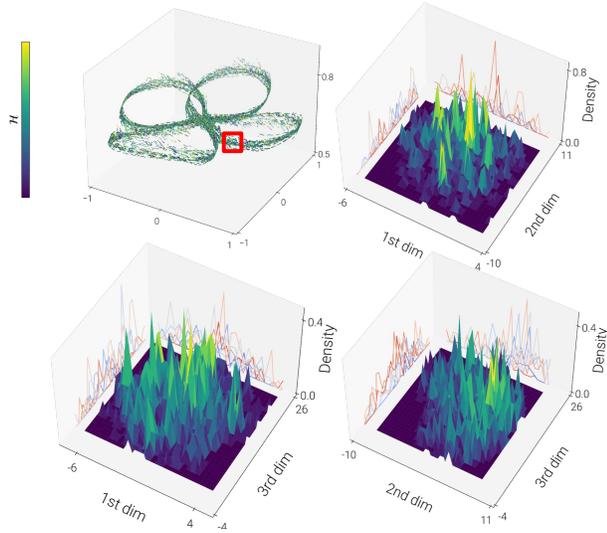


Fig. 4. *Top left*: Position of the quadrotor during real-world flights with the uncertainties at each timestep \mathcal{H}_t highlighted. The following plots illustrate the empirical density of \mathcal{H} observed in the **red box** along two dimensions. *Top right*: Density along first and second dimensions, *Bottom left*: Density along the first and third dimension, and *Bottom right*: Density along the second and third dimensions. Even under ideal conditions with no external disturbances, \mathcal{H} observed in a given flight path segment (expected to have similar control inputs and sensor feedback) exhibits a multimodal nature.

Quadrotor dynamics are inherently multimodal, driven by a combination of environmental and mechanical factors that introduce significant uncertainties. Turbulent wind

gusts, temperature-induced variations in air density, and aerodynamic effects such as propeller-induced flows and ground effects result in nonlinearities. Mechanical factors, including motor degradation, IMU sensor noise, and control delays, further complicate the structure of \mathcal{H} . Many of these parameters—such as air density, humidity, and downwash effects—are difficult to measure with onboard sensors, and their variations due to altitude and temperature fluctuations amplify uncertainties during flight. Repeated flights along an infinity-shaped trajectory with a fixed PID controller (Figure 4, top left) reveal varying levels of \mathcal{H} even at identical flight segments. The remaining panels in Figure 4 confirm the multimodal structure of \mathcal{H} in real-world operating regimes. The empirical density of \mathcal{H} for a specific segment of the trajectory (highlighted in **red**) reveals distinct modes, indicating variability in the dynamics even without external disturbances (Refer to [29, Appendix A] for quantitative analysis and further discussion).

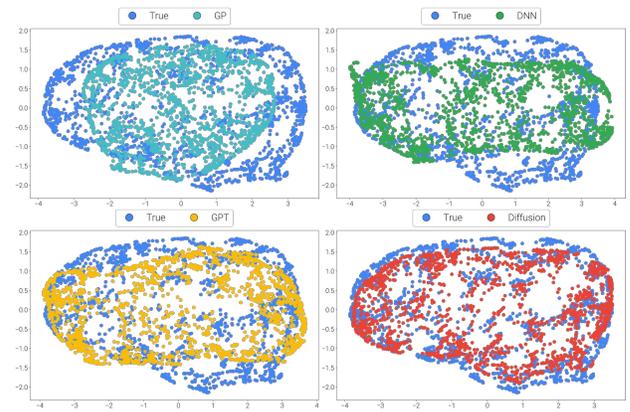


Fig. 5. t-SNE plots of \mathcal{H} observed from real-world flight data (**red box** in Figure 4) and \mathcal{H} obtained from the baselines and the proposed Diffusion model. While GPT fairly estimates the uncertainties in the dynamics, the proposed diffusion model accurately captures the underlying distribution.

If \mathcal{H}_t is modeled as independent multimodal distributions, such as in Gaussian Mixture Models, successive predictions may originate from different modes, leading to temporal inconsistencies in the dynamics [28]. To address this issue, we leverage the ability of DDPMs to capture the distribution of \mathcal{H} . Unlike traditional methods, DDPMs leverage Stochastic Langevin Dynamics [30] in the denoising process, allowing them to capture multimodal distributions that naturally reflect the inherent stochasticity of the system. Building on the success of observation-conditioned visuomotor policies [28], we use the DDPM to approximate the conditional distribution $p(\mathcal{H}|\zeta, u)$, instead of the joint distribution $p(\mathcal{H}, \zeta, u)$ typically done in planning [17]. This allows us to better capture the uncertainties by conditioning on sensory inputs ζ and control inputs u . To achieve this, we modify the reverse process of DDPM as follows:

$$\mathcal{H}^{k-1} = \frac{1}{\sqrt{\alpha^k}} \left(\mathcal{H}^k - \frac{1 - \alpha^k}{\sqrt{1 - \bar{\alpha}^k}} \epsilon_\theta(\zeta, u, \mathcal{H}^k, k) + \sqrt{\beta^k} \mathbf{z} \right), \quad (5)$$

where $\alpha^k := 1 - \beta^k$, $\bar{\alpha}^k := \prod_{i=1}^k \alpha^i$ and $\mathbf{z} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$. With an offline dataset $\mathcal{D} = \{\mathcal{S}\}_{i=1}^M$ and the denoising process (5), the simplified objective (1) becomes:

$$\mathcal{J}(\theta) = \mathbb{E}_{\substack{k \sim \mathcal{U}(1, K), \epsilon \sim \mathcal{N}(\mathbf{0}, \mathbf{I}) \\ (\zeta, u, \mathcal{H}^0) \sim \mathcal{D}}} \left[\|\epsilon - \epsilon_\theta(\zeta, u, \mathcal{H}^k, k)\|^2 \right]. \quad (6)$$

We now discuss the deployment of the trained diffusion model in real-world flight scenarios. The process begins by specifying the desired waypoints ζ_d and collecting real-time sensory data ζ for trajectory tracking. Using the trained model, we generate a conditional sample $\hat{\mathcal{H}}$ via (5), where the covariance follows the cosine schedule proposed in [31]. The uncertainty for the first timestep $\hat{\mathcal{H}}_1$ from the resulting estimate is then utilized to compute the control inputs u . This procedure is iteratively applied within a receding-horizon control framework, as depicted in Figure 1, ensuring ongoing adaptation to the dynamic flight conditions. Figure 5 illustrates the effectiveness of the proposed diffusion model in capturing the intricate structure of \mathcal{H} .

D. Diffusion Model-based Adaptive Control (DM-AC)

Given the estimate $\hat{\mathcal{H}}$ obtained from the proposed diffusion model, we now proceed to design a closed-loop feedback controller to ensure that the position tracking error $e_p(t) \triangleq p(t) - p_d(t)$ converges to zero, ideally at an exponential rate. This can be achieved through the design of a sliding variable s , which defines a manifold that guarantees the desired convergence rate of the tracking error:

$$s(t) = \dot{e}_p(t) + \Phi e_p(t), \quad (7)$$

where Φ is a user-defined positive definite gain matrix. The adaptive controller can be derived as:

$$u(t) = -\Lambda s + \bar{m}(\ddot{p}_d - \Phi \dot{e}_p) + \hat{\mathcal{H}} - \hat{\sigma}(t)(s/\|s\|), \quad (8a)$$

$$\dot{\hat{\sigma}}(t) = \|s\| - \nu \hat{\sigma}(t), \quad \hat{\sigma}(0) > 0, \quad (8b)$$

where Λ is a positive definite user-defined gain matrix, $\nu \in \mathbb{R}^+$ are user-defined scalars, and $\hat{\sigma}(t)$ is adaptive gain. This not only ensures the dynamic compensation through $\hat{\mathcal{H}}$ but also accounts for estimation errors in $\hat{\mathcal{H}}$ resulting from unknown external disturbances. Theorem II.3 ensures the closed-loop stability of the system using DM-AC (refer to [29, Appendix B] for proof).

Assumption II.1. *The desired states along the position trajectory p_d , \dot{p}_d , and \ddot{p}_d are bounded. Also, the external disturbances are bounded.*

Assumption II.2. *The learning error of $\hat{\mathcal{H}}$ over the compact sets $\zeta \in \mathcal{Z}$, $u \in \mathcal{U}$ is upper bounded by $\sigma_m = \sup_{\zeta, u} \|\sigma\|$, where $\sigma = \mathcal{H} - \hat{\mathcal{H}}$.*

Theorem II.3. *Under Assumptions II.1, II.2 the closed-loop trajectories of (3) with control law (8a) along with the adaptive law (8b) are Uniformly Ultimately Bounded (UUB) (Definition 4.6 in [32]).*

Remark 1. *While Assumption II.2 is necessary for ensuring system stability (Theorem II.3), it is well-founded and aligns*

with recent findings for Score-based Generative Models (SGM) [25]. Works like [33]–[35] show that SGMs can achieve convergence to the data distribution by bounding the divergence between the learned and true distributions. This suggests that Assumption II.2 is satisfied when the training data is sufficiently representative of the underlying distribution of \mathcal{H} . Moreover, σ_m is unknown for control design of (8).

Algorithm 1 Quadrotor Trajectory Tracking with DM-AC

Require: Noise model ϵ_θ , covariances Σ^k , gains $\bar{m}, \Lambda, \Phi, \nu$

Initialise: $u = 0, \hat{\sigma}(0)$

```

1: while not done do
2:   Get desired way-point  $\zeta_d = (p_d, \dot{p}_d, \ddot{p}_d)$ 
3:   Observe sensor feedback  $\zeta$ , initialise  $\mathcal{H}^T \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ 
4:   for  $k=T, \dots, 1$  do
5:      $(\mu^{k-1}, \Sigma^k) \leftarrow \text{Denoise}(\mathcal{H}^k, \epsilon_\theta(\zeta, u, \mathcal{H}^k, k))$ 
6:      $\mathcal{H}^{k-1} \sim \mathcal{N}(\mu^{k-1}, (1 - \beta)\Sigma^{k-1})$ 
7:   end for
8:   Derive  $e_p, \dot{e}_p$  and calculate  $s$  as in (7)
9:   Calculate  $\hat{\sigma}$  using (8b)
10:  Update  $u$  using (8a)
11: end while

```

Algorithm 1 details the use of the diffusion sample $\hat{\mathcal{H}}$ with the adaptive controller (8) for trajectory tracking¹. The desired total thrust and the desired attitude can be calculated as per [36].

III. EXPERIMENTS

The primary objective of our experiments is to rigorously evaluate the efficacy of the proposed methodology in meeting the performance expectations of a data-driven trajectory tracking framework. Specifically, we aim to address the following key questions: **(1)** How accurately can the framework track complex flight paths that were not encountered during training? **(2)** How well does the framework generalize to varying dynamic conditions and external disturbances beyond the training distribution? We conclude by studying practical runtime considerations of the framework.

Simulation Setup: We utilize *Iris* quadrotor, deployed on the open-source *PX4-Autopilot* flight control software, integrated within the Gazebo simulation environment.

Hardware Setup: The experimental setup includes an 8-camera OptiTrack motion capture system and a quadrotor (approximately 1.4 kg) featuring DYS D4215 - 650 KV brushless motors, powered by a 4S (120 C) LiPo battery, and equipped with 10-inch propellers on an F-450 frame, as shown in Figure 1 (bottom-right). The quadrotor is equipped with a Pixhawk flight controller running PX4, an onboard Jetson Xavier NX computer running the MAVROS package, and a host computer with an NVIDIA GeForce RTX 4080 for inference. Sensor data is transmitted from the onboard computer to the host computer via WiFi, where it is processed,

¹An alternative control law is presented in [29, Appendix C], demonstrating the flexibility of the framework.

and control inputs are sent back to the onboard computer at a frequency of 30 Hz. The hardware configuration reflects a standard setup typical of consumer drone systems. The PX4 is operated in offboard mode, receiving desired thrust and attitude commands from the position control loop (8). The built-in PX4 multicopter attitude controller (linear PID controller based on quaternion error) was executed at the default rate. Desired trajectories were generated using the open-source ROS package ‘*mav_trajectory_generation*’ [37].

Implementation: We construct an offline dataset of sequences \mathcal{S} by computing \mathcal{H} from (3) and train the diffusion model to minimize (6). The noise predictor ϵ_θ is parameterized by a temporal U-Net [17], which uses three repeated residual blocks, each with two temporal convolutions, followed by group normalization [38] and Mish activation [39]. Timestep and condition embeddings (16-dimensional vectors) are produced by separate 2-layer MLPs with 64 hidden units and Mish activation. We train ϵ_θ using the Adam optimizer [40] with a learning rate 2×10^{-4} and a batch size of 256 for 50,000 steps, using $K = 20$ diffusion timesteps. The adaptive controller gains are set as: $\Phi = \text{diag}\{1.5, 1.5, 1.2\}$, $\Lambda = \text{diag}\{2.0, 2.0, 4.0\}$, with $\bar{m} = 1$, $\hat{\sigma}(0) = 0.1$, and $\nu = 2.0$.

Baselines: For a baseline comparison, we implement a first-principles-based \mathcal{L}_1 -Adaptive controller [41]², which is widely recognized for its proven stability and performance in adaptive control tasks. To demonstrate the effectiveness of the proposed diffusion-based framework, we compare it against strong, established data-driven model-learning approaches, including GPs [3]³ and DNNs [12], which have shown success in learning quadrotor dynamics, and autoregressive models such as GPT-2 [42]⁴, commonly used for sequence prediction tasks. For a fair tracking comparison, all models were trained on the same offline dataset to predict $\hat{\mathcal{H}}$ and were used with the control law (8). Results are reported for 10 individual trials.

A. Generalization to Unseen Complex Trajectories

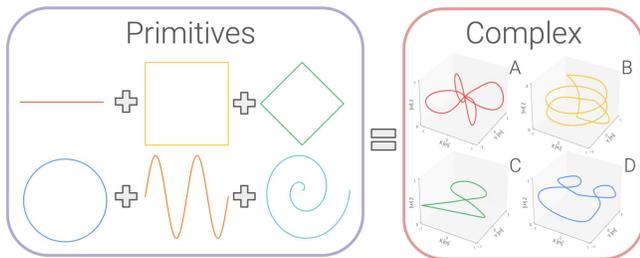


Fig. 6. *Left:* Primitive trajectories, such as straight lines, squares, kites, circles, sinusoidal waves, and spirals, were executed in all three planes (xy, yz, xz) to capture a variety of motion patterns. *Right:* Complex trajectories include a combination of sharp turns, loops, and simultaneous motion in all three (x, y, z) coordinates, illustrating challenging real-world flight scenarios.

To evaluate the generalization capability of the proposed framework to previously unseen environments, we trained the

²<https://github.com/sigma-pi/L1Quad>

³https://github.com/uzh-rpg/data_driven_mpc

⁴<https://github.com/kzll/decision-transformer>

models on single demonstrations of basic *primitive trajectories* (Figure 6) with a mean velocity of 0.4 m/s and tested on more *complex trajectories* in a simulated environment. Our evaluation focuses on both predictive accuracy and trajectory tracking performance. For predictive accuracy, we

TABLE I
PREDICTION AND TRACKING COMPARISON IN SIMULATION

RMSE	Model Prediction Error (N)				Trajectory Tracking Error (m)			
	A	B	C	D	A	B	C	D
GP-AC	0.097	0.108	0.055	0.073	0.047	0.052	0.039	0.030
DNN-AC	0.091	0.101	0.046	0.062	0.032	0.047	0.031	0.025
GPT-AC	0.088	0.089	0.044	0.048	0.031	0.038	0.028	0.022
DM-AC	0.073	0.081	0.041	0.051	0.022	0.021	0.019	0.023

*{ }-AC refers to { } based adaptive control.

tasked the quadrotor to follow the target trajectory using a vanilla PID controller while calculating \mathcal{H} at each timestep from (3). We compare the predictive capabilities of the four baselines—GP, DNN, GPT, and our proposed diffusion model—using the features $\{\zeta, u\}$ at each timestep. We then assess trajectory tracking performance by implementing the control law (8) and computing the root mean square (RMS) position errors for each method, as shown in Table I. The results demonstrate that the diffusion model consistently achieves superior predictive accuracy and tracking performance, exhibiting lower RMS errors compared to the other baselines. While GPT performs comparably to the diffusion model on one of the *complex trajectories* (Complex D), the diffusion model outperforms all baselines across the remaining cases.

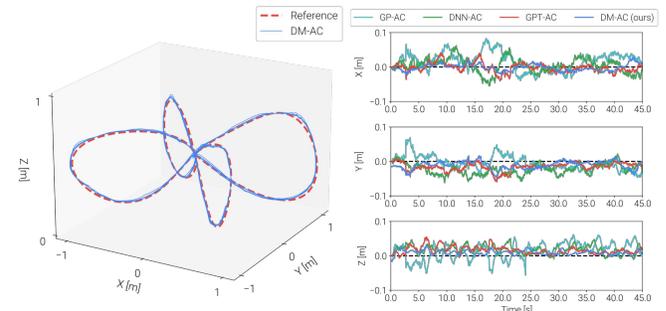


Fig. 7. *Left:* Comparison between reference and tracked trajectory in real-world flight for DM-AC. *Right:* Tracking error comparison on Complex ‘A’ in real-world flight. RMS tracking errors for GP-AC, DNN-AC, GPT-AC, and DM-AC are 0.075 m, 0.071 m, 0.067 m and 0.054 m, respectively.

Given that aerodynamic effects are more accurately captured in real-world flight than in simulation, we further validate the framework using real-world data. Specifically, we evaluate the performance of DM-AC on the Complex A trajectory (Figure 6), using primitive demonstrations recorded at a velocity of 0.4m/s for training. As shown in Figure 7, DM-AC achieves the lowest tracking error. This highlights the ability of DM-AC to track complex, unseen trajectories without requiring extensive data collection for each specific trajectory. This flexibility is driven by the learned conditional distribution $p(\mathcal{H}|\zeta, u)$, which enables DM-AC to effectively generalize across a wide range of challenging real-world conditions and scenarios.

B. Adaptation against various Payload, Velocities and Wind

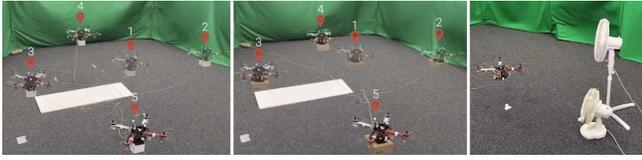


Fig. 8. *Left and Centre*: Sequential snapshots of the quadrotor with payloads 0.4 kg and 0.8 kg, respectively. *Right*: Quadrotor in the presence of external wind generated by two fans.

To assess the adaptability of the proposed framework, we test on a complex double infinity-shaped trajectory (where the first major loop is aligned along the x -axis and the second along the y -axis, as depicted in the top left of Figure 4) under varying payloads and wind disturbances. The evaluation is conducted under the following conditions: (a) tracking at velocities of 0.35 m/s and 0.5 m/s with a payload of 0.4 kg, (b) tracking at the same velocities with a payload of 0.8 kg, and (c) hovering at (0, 0, 0.8) m under external wind disturbances (Figure 8). The diffusion model was trained on *single real-world demonstrations of primitive trajectories* (left Figure 8) with payloads of {0.2, 0.6} kg and velocities of {0.2, 0.4} m/s. Notably, the payload configuration of 0.4 kg and velocity of 0.35 m/s are within the range of the training data but not explicitly used for training. In contrast, the 0.8 kg payload and 0.5 m/s velocity represent extrapolated scenarios that fall outside the training distribution.

TABLE II
PERFORMANCE COMPARISON IN REAL WORLD

RMSE (m)	Payload = 0.4 kg		Payload = 0.8 kg		External Wind
	$v = 0.35$ m/s	$v = 0.5$ m/s	$v = 0.35$ m/s	$v = 0.5$ m/s	
\mathcal{L}_1 -Adaptive	0.088	0.101	0.105	0.112	0.198
GP-AC	0.086	0.098	0.102	0.110	0.192
DNN-AC	0.082	0.095	0.098	0.107	0.184
GPT-AC	0.075	0.086	0.089	0.102	0.177
DM-AC	0.069	0.074	0.077	0.089	0.162

*{·}-AC refers to {·} based adaptive control.

Table II presents a comparison of the tracking performance across different model-based methods. The results clearly show that the model-based framework, which integrates adaptive control with a learned dynamic model, significantly outperforms the \mathcal{L}_1 -adaptive controller, particularly in scenarios with varying payloads and external wind disturbances. The combined approach enables effective short-term adjustments and long-term compensation for dynamic variations, ensuring both robustness and accuracy. Additionally, the GP and DNN-based methods are less effective at handling the real-world multimodal uncertainties present in these scenarios, resulting in reduced tracking accuracy. While GPT achieves reasonable performance, its autoregressive nature introduces cascading errors, leading to slightly higher tracking errors. In contrast, DM-AC's ability to generate accurate multistep predictions enhances its performance, demonstrating resilience to variations in dynamic parameters and external disturbances.

C. Runtime performance with multi-step predictions

Sampling from diffusion models generally entails considerable computational expense due to the iterative nature of the denoising process during inference. In our framework, \mathcal{H} is predicted over a finite time horizon H (horizon length), which permits the sequential use of multistep predictions to potentially decrease the frequency of sampling from the diffusion model. However, employing predictions across multiple timesteps to calculate the control inputs u may adversely affect tracking performance. To investigate the impact of computational budget on performance, we tasked the quadrotor to stabilize at a fixed coordinate of (0.0, 0.0, 0.8) m for 60 s using DM-AC while increasing the number of predictions used from a sample. Figure 9 illustrates the trade-off between performance and Execution Time per time Elapsed over the prediction horizon (ETE). The ETE is defined as the ratio of the time spent evaluating the diffusion model to the time elapsed between successive queries. Our findings reveal that stabilization accuracy is maintained with no crashes when the number of predictions used from a sample is low. However, as the number of predictions is increased to 32 and 64, stability deteriorates resulting in crash rates of 20% and 60%, respectively. This degradation in performance can be attributed to the reliance on outdated observations of states and control inputs for estimating \mathcal{H} .

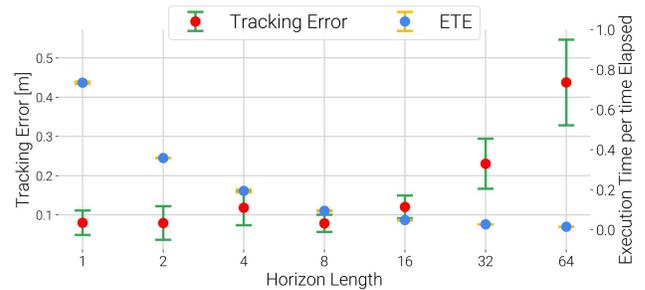


Fig. 9. Hovering accuracy evaluation and execution time for 60 seconds with different horizon lengths H .

IV. CONCLUSION AND FUTURE WORK

In this work, we introduced *DroneDiffusion*, a novel framework that effectively leverages diffusion models to capture the stochastic and multimodal nature of quadrotor dynamics. Through training on primitive demonstrations, *DroneDiffusion* exhibits strong generalization capabilities, accurately tracking complex flight paths, including scenarios not encountered during training. The framework's robustness was further validated on real hardware, demonstrating its adaptability to a wide range of payload and velocity variations, as well as its resilience against wind disturbances. These findings establish *DroneDiffusion* as a promising solution for real-world deployment in dynamic and uncertain environments. Future work can leverage recent advances in efficient sampling techniques [43], [44] to enable on-board inference. Additionally, the potential for extending the framework to agile maneuvering and attitude control is a compelling direction.

REFERENCES

- [1] H. Yao and X. Liang, "Autonomous exploration under canopy for forest investigation using lidar and quadrotor," *IEEE Transactions on Geoscience and Remote Sensing*, 2024.
- [2] M. Idrissi, M. Salami, and F. Annaz, "A review of quadrotor unmanned aerial vehicles: applications, architectural design and control algorithms," *Journal of Intelligent & Robotic Systems*, vol. 104, no. 2, p. 22, 2022.
- [3] G. Torrente, E. Kaufmann, P. Föhn, and D. Scaramuzza, "Data-driven mpc for quadrotors," *IEEE Robotics and Automation Letters*, vol. 6, no. 2, pp. 3769–3776, 2021.
- [4] M. Labbadi and M. Cherkaoui, "Robust adaptive nonsingular fast terminal sliding-mode tracking control for an uncertain quadrotor uav subjected to disturbances," *ISA transactions*, vol. 99, pp. 290–304, 2020.
- [5] W. Xie, D. Cabecinhas, R. Cunha, and C. Silvestre, "Adaptive backstepping control of a quadcopter with uncertain vehicle mass, moment of inertia, and disturbances," *IEEE Transactions on Industrial Electronics*, vol. 69, no. 1, pp. 549–559, 2021.
- [6] F. Crocetti, J. Mao, A. Saviolo, G. Costante, and G. Loianno, "Gapt: Gaussian process toolkit for online regression with application to learning quadrotor dynamics," in *2023 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2023, pp. 11 308–11 314.
- [7] N. Schmid, J. Gruner, H. S. Abbas, and P. Rostalski, "A real-time gp based mpc for quadcopters with unknown disturbances," in *2022 American Control Conference (ACC)*. IEEE, 2022, pp. 2051–2056.
- [8] L. Wang, E. A. Theodorou, and M. Egerstedt, "Safe learning of quadrotor dynamics using barrier certificates," in *2018 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2018, pp. 2460–2465.
- [9] A. Punjani and P. Abbeel, "Deep learning helicopter dynamics models," in *2015 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2015, pp. 3223–3230.
- [10] S. Bansal, A. K. Akametalu, F. J. Jiang, F. Laine, and C. J. Tomlin, "Learning quadrotor dynamics using neural network for flight control," in *2016 IEEE 55th Conference on Decision and Control (CDC)*. IEEE, 2016, pp. 4653–4660.
- [11] Q. Li, J. Qian, Z. Zhu, X. Bao, M. K. Helwa, and A. P. Schoellig, "Deep neural networks for improved, impromptu trajectory tracking of quadrotors," in *2017 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2017, pp. 5183–5189.
- [12] G. Shi, X. Shi, M. O'Connell, R. Yu, K. Azizzadenesheli, A. Anandkumar, Y. Yue, and S.-J. Chung, "Neural lander: Stable drone landing control using learned dynamics," in *2019 international conference on robotics and automation (icra)*. IEEE, 2019, pp. 9784–9790.
- [13] N. Mohajerin, M. Mozifian, and S. Waslander, "Deep learning a quadrotor dynamic model for multi-step prediction," in *2018 IEEE international conference on robotics and automation (ICRA)*. IEEE, 2018, pp. 2454–2459.
- [14] N. Mohajerin and S. L. Waslander, "Multistep prediction of dynamic systems with recurrent neural networks," *IEEE transactions on neural networks and learning systems*, vol. 30, no. 11, pp. 3370–3383, 2019.
- [15] A. Radford, J. Wu, R. Child, D. Luan, D. Amodei, I. Sutskever, et al., "Language models are unsupervised multitask learners," *OpenAI blog*, vol. 1, no. 8, p. 9, 2019.
- [16] A. Venkatraman, M. Hebert, and J. Bagnell, "Improving multi-step prediction of learned time series models," in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 29, no. 1, 2015.
- [17] M. Janner, Y. Du, J. Tenenbaum, and S. Levine, "Planning with diffusion for flexible behavior synthesis," in *International Conference on Machine Learning*, 2022.
- [18] S. Sun, C. C. de Visser, and Q. Chu, "Quadrotor gray-box model identification from high-speed flight data," *Journal of Aircraft*, vol. 56, no. 2, pp. 645–661, 2019.
- [19] L. Bauersfeld, E. Kaufmann, P. Foehn, S. Sun, and D. Scaramuzza, "Neurobem: Hybrid aerodynamic quadrotor model," *ArXiv*, vol. abs/2106.08015, 2021. [Online]. Available: <https://api.semanticscholar.org/CorpusID:235435822>
- [20] L. Bauersfeld, K. Muller, D. Ziegler, F. Coletti, and D. Scaramuzza, "Robotics meets fluid dynamics: A characterization of the induced airflow around a quadrotor," *arXiv preprint arXiv:2403.13321*, 2024.
- [21] A. Saviolo, J. Frey, A. Rathod, M. Diehl, and G. Loianno, "Active learning of discrete-time dynamics for uncertainty-aware model predictive control," *IEEE Transactions on Robotics*, 2023.
- [22] A. Saviolo and G. Loianno, "Learning quadrotor dynamics for precise, safe, and agile flight control," *Annual Reviews in Control*, 2023.
- [23] J. Sohl-Dickstein, E. Weiss, N. Maheswaranathan, and S. Ganguli, "Deep unsupervised learning using nonequilibrium thermodynamics," in *International conference on machine learning*. PMLR, 2015, pp. 2256–2265.
- [24] J. Ho, A. Jain, and P. Abbeel, "Denoising diffusion probabilistic models," *Advances in neural information processing systems*, vol. 33, pp. 6840–6851, 2020.
- [25] Y. Song, J. Sohl-Dickstein, D. P. Kingma, A. Kumar, S. Ermon, and B. Poole, "Score-based generative modeling through stochastic differential equations," in *International Conference on Learning Representations*, 2020.
- [26] P. Dhariwal and A. Nichol, "Diffusion models beat gans on image synthesis," *Advances in neural information processing systems*, vol. 34, pp. 8780–8794, 2021.
- [27] A. Ajay, Y. Du, A. Gupta, J. B. Tenenbaum, T. S. Jaakkola, and P. Agrawal, "Is conditional generative modeling all you need for decision making?," in *The Eleventh International Conference on Learning Representations*, 2023. [Online]. Available: <https://openreview.net/forum?id=sP1fo2K9DFG>
- [28] C. Chi, S. Feng, Y. Du, Z. Xu, E. Cousineau, B. Burchfiel, and S. Song, "Diffusion policy: Visuomotor policy learning via action diffusion," in *Proceedings of Robotics: Science and Systems (RSS)*, 2023.
- [29] A. Das, R. D. Yadav, S. Sun, M. Sun, S. Kaski, and W. Pan, "Dronediffusion: Robust quadrotor dynamics learning with diffusion models," *arXiv preprint arXiv:2409.11292*, 2024.
- [30] M. Welling and Y. W. Teh, "Bayesian learning via stochastic gradient langevin dynamics," in *Proceedings of the 28th international conference on machine learning*. Citeseer, 2011, pp. 681–688.
- [31] A. Q. Nichol and P. Dhariwal, "Improved denoising diffusion probabilistic models," in *International conference on machine learning*. PMLR, 2021, pp. 8162–8171.
- [32] H. K. Khalil, *Nonlinear control*. Pearson New York, 2015, vol. 406.
- [33] H. Lee, J. Lu, and Y. Tan, "Convergence for score-based generative modeling with polynomial complexity," *Advances in Neural Information Processing Systems*, vol. 35, pp. 22 870–22 882, 2022.
- [34] V. De Bortoli, "Convergence of denoising diffusion models under the manifold hypothesis," *Transactions on Machine Learning Research*, 2022.
- [35] S. Chen, S. Chewi, J. Li, Y. Li, A. Salim, and A. R. Zhang, "Sampling is as easy as learning the score: theory for diffusion models with minimal data assumptions," in *International Conference on Learning Representations*, 2023.
- [36] D. Mellinger and V. Kumar, "Minimum snap trajectory generation and control for quadrotors," in *2011 IEEE international conference on robotics and automation*. IEEE, 2011, pp. 2520–2525.
- [37] A. S. L. ETH Zurich. mav trajectory generation. [Online]. Available: https://github.com/ethz-asl/mav_trajectory_generation
- [38] Y. Wu and K. He, "Group normalization," in *Proceedings of the European conference on computer vision (ECCV)*, 2018, pp. 3–19.
- [39] D. Misra, "Mish: A self regularized non-monotonic activation function," in *British Machine Vision Conference*, 2020. [Online]. Available: <https://api.semanticscholar.org/CorpusID:221113156>
- [40] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," *arXiv preprint arXiv:1412.6980*, 2014.
- [41] Z. Wu, S. Cheng, P. Zhao, A. Gahlawat, K. A. Ackerman, A. Lakshmanan, C. Yang, J. Yu, and N. Hovakimyan, "L1 quad: L1 adaptive augmentation of geometric control for agile quadrotors with performance guarantees," *arXiv preprint arXiv:2302.07208*, 2023.
- [42] L. Chen, K. Lu, A. Rajeswaran, K. Lee, A. Grover, M. Laskin, P. Abbeel, A. Srinivas, and I. Mordatch, "Decision transformer: Reinforcement learning via sequence modeling," *Advances in neural information processing systems*, vol. 34, pp. 15 084–15 097, 2021.
- [43] K. Zheng, C. Lu, J. Chen, and J. Zhu, "Dpm-solver-v3: Improved diffusion ode solver with empirical model statistics," *Advances in Neural Information Processing Systems*, vol. 36, 2024.
- [44] Y. Song, P. Dhariwal, M. Chen, and I. Sutskever, "Consistency models," in *International Conference on Machine Learning*. PMLR, 2023, pp. 32 211–32 252.