Physics-Informed Probabilistic Learning of Low-Altitude Urban Wind from Onboard Motion Data

Minghao Chen, Antonina Kosikova, Marco G. Giometto, Andrew W. Smyth *
Department of Civil Engineering and Engineering Mechanics
Columbia University
New York, NY 10025
minghao.c, ak4967, mg3929, aws16@columbia.edu

Abstract

Urban wind fields at low altitude are highly variable and safety-critical for urban air mobility and drone operations. This paper presents a physics-informed probabilistic machine learning approach to estimate 3D local winds: a quadrotor dynamics prior coupled with a learned discrepancy, inferred from onboard IMU data only. The latent wind vector and its uncertainty can be estimated in near-real time, enabling risk-aware planning in cities. Unlike existing methods that rely on dedicated sensors or complex models with high computational demands, our approach uses only motion data from standard onboard IMU sensors. Our results demonstrate that the method can estimate wind vectors with a root mean square error (RMSE) of less than 12% (of speed range) under various simulated realistic wind conditions in both hovering and cruising cases, while providing uncertainty quantification of the estimates. The proposed framework offers an easy-to-implement low-altitude onboard wind estimator to support drone operations in urban environments.

1 Introduction

With the recent development of unmanned aerial vehicles (UAVs), many stakeholders and researchers are preparing for Urban Air Mobility (UAM) as an important part of future transportation. Technological improvements in electric propulsion, high-capacity batteries and air traffic control, along with increasing aerial autonomy, are positioning UAM to revolutionize urban transportation and commercial aviation. Companies have begun deploying air taxis and drone delivery services (e.g., Amazon [9] and Google [11]), while administrative agencies such as FAA and NASA are working closely to make new procedures, requirements, and protocols to enable higher operation densities [2].

Nevertheless, there are still risks that need to be addressed before new entrants can safely access the airspace. Low altitude turbulence, particularly the complex winds within and immediately above the urban canopy, can significantly impact the safety and efficiency of vehicles en route to their destinations ([1], [3]). These challenging wind conditions can lead to service delays, delivery cancellations, and even loss of the vehicle control. As a result, customers and city officials express heightened concerns about the safety, reliability, and robustness of UAM compared to current ground transportation. Traditional meteorological observations often lack the precision needed for estimating local wind conditions. Computational fluid dynamics (CFD) simulations and spatially-distributed wind sensors (anemometers) have been explored as a possible pathway, but usually are costly and, to date, are unable to be updated in real time to adapt to time-varying urban wind fields. Accessible and affordable real-time wind information can significantly aid UAV operations, supporting energy-efficient and safe path planning for the full-scale adoption of UAM. This data gap highlights the need

*Corresponding author

39th Conference on Neural Information Processing Systems (NeurIPS 2025) Workshop: Urban AI.

for adaptive, data-driven models that can learn from real-time observations and operate efficiently on the resource-constrained computational platforms typical of UAVs.

To address these challenges, we study UAV-as-a-sensor to realize fast mobile wind sensing. We propose a probabilistic learning approach that fuses a simplified quadrotor dynamics model with onboard data to infer wind, treating the wind as a latent state with priors and learning noise/discrepancy statistics. The method infers wind velocity while also computing uncertainties associated with the estimates, ensuring higher safety margins and robustness against unpredictable environmental factors.

2 Methodology

Our primary goal is to infer the latent variable, the wind vector $\boldsymbol{w}$, from a time series of noisy IMU acceleration measurements $\mathbf{Y}$, which can be formulated as a probabilistic inference problem.

Generative Model of UAV Motion. We model the observed acceleration vector $\boldsymbol{y}_i$ at time step as a realization from a generative process defined by:

$$\boldsymbol{y}_i = \boldsymbol{f}(\boldsymbol{w}_i, \mathbf{z}_{q,i}, \mathbf{U}_i) + \boldsymbol{\epsilon}_i$$

where $\boldsymbol{f}(\cdot)$ is a deterministic 6-DoF quadcopter dynamics model in terms of vehicle states $z_{q,i}$ (e.g., orientation and quaternion), controller output $\mathbf{U_i}$, and wind vector. $\boldsymbol{\epsilon}_i$ is a stochastic error term. The simplified dynamics function $\boldsymbol{f}$ is defined as [5]:

$$\boldsymbol{f} = \frac{1}{m_B}\left(\mathbf{D}(\boldsymbol{w}, \boldsymbol{v}) + R(\mathbf{z_{q,i}})|\mathbf{U}| + \mathbf{G}\right)$$

where $R(\cdot)$ is the rotation matrix in terms of states, $\mathbf{D}(\cdot)$ is the aerodynamic drag force which non-linearly depends on the relative velocity $(\boldsymbol{w} - \boldsymbol{v})$, bridging the wind with acceleration. The error term is modeled as zero-mean Gaussian prior, $\boldsymbol{\epsilon}_i \sim \mathcal{N}(\mathbf{0}, \boldsymbol{\Sigma}_\epsilon)$, where $\boldsymbol{\Sigma}_\epsilon = \mathrm{diag}(\sigma^2_{\epsilon,x}, \sigma^2_{\epsilon,y}, \sigma^2_{\epsilon,z})$. This is a flexible and standard assumption in probabilistic modeling for handling structured noise and model discrepancy.

Bayesian Inference. To compute the posterior $p(\boldsymbol{w} \mid \mathbf{Y}) \propto p(\mathbf{Y} \mid \boldsymbol{w})p(\boldsymbol{w})$ and perform the inference, the likelihood function is defined as follows. To resolve the curse of the non-linearity of $\boldsymbol{f}(\cdot)$, we assume the wind within a data segment is constant and follows a Gaussian distribution, $\boldsymbol{w} \sim \mathcal{N}(\boldsymbol{\mu}_w, \boldsymbol{\Sigma}_w)$. Previous studies ([13] and [15]) have shown mean wind force remains constant within 1-5 seconds, implying it could be invariant within a segment of samples. We then linearize

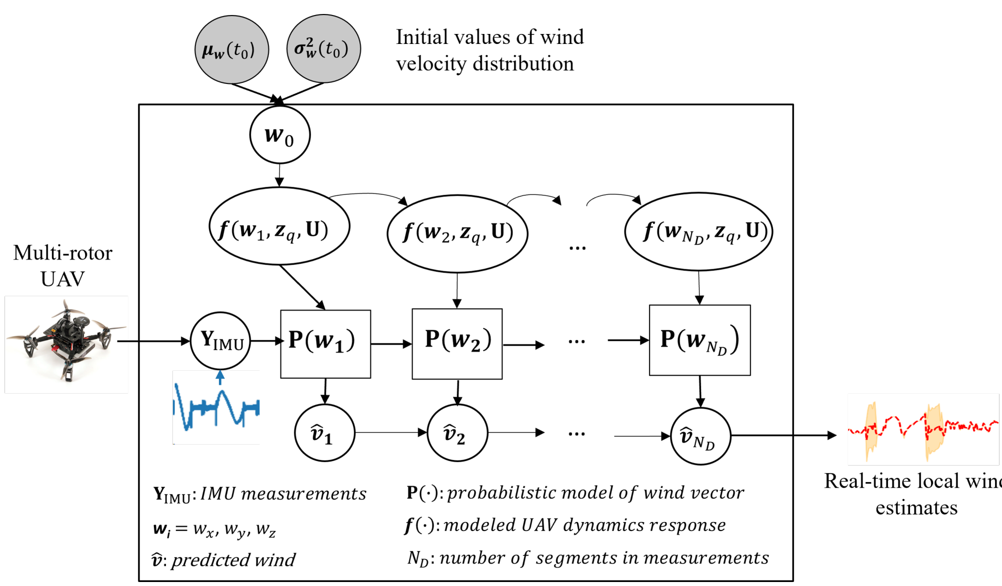

Figure 1: Overview of the proposed framework.

$\boldsymbol{f}(\cdot)$ via a first-order Taylor series expansion around $\boldsymbol{\mu}_w$:

$$\boldsymbol{f}(\boldsymbol{w}) \approx \boldsymbol{f}(\boldsymbol{\mu}_w) + \mathbf{J}(\boldsymbol{w} - \boldsymbol{\mu}_w) \quad \text{where} \quad \mathbf{J} = \left.\frac{\partial \boldsymbol{f}}{\partial \boldsymbol{w}}\right|_{\boldsymbol{w}=\boldsymbol{\mu}_w}$$

This approximation allows us to marginalize the latent variable $\boldsymbol{w}$, yielding a tractable Gaussian likelihood for the observations $\mathbf{Y}$ given the wind distribution parameters:

$$p(\mathbf{Y} \mid \boldsymbol{\mu}_w, \boldsymbol{\Sigma}_w) \sim \mathcal{N}\left(\boldsymbol{f}(\boldsymbol{\mu}_w), \mathbf{J}\boldsymbol{\Sigma}_w\mathbf{J}^T + \boldsymbol{\Sigma}_\epsilon\right)$$

The covariance term fuses the propagated uncertainty from the wind estimate ($\mathbf{J}\boldsymbol{\Sigma}_w\mathbf{J}^T$) with the measurement noise ($\boldsymbol{\Sigma}_\epsilon$).

Posterior Optimization. The posterior parameters $\boldsymbol{\nu} = \{\boldsymbol{\mu}_w, \boldsymbol{\Sigma}_w\}$ are inferred by minimizing the negative log-likelihood for each data segment. To regularize the problem, we incorporate prior knowledge by constraining the mean wind magnitude based on meteorological data:

$$\hat{\boldsymbol{\nu}} = \text{argmin}_{\boldsymbol{\nu}} \left(-\log p(\mathbf{Y} \mid \boldsymbol{\nu})\right) \quad \text{s.t.} \quad \|\boldsymbol{\mu}_w\| \in [\alpha, \beta]$$

This constrained optimization yields a point estimate of the wind vector, $\hat{\boldsymbol{\mu}}_w$, and its associated uncertainty, $\hat{\boldsymbol{\Sigma}}_w$. α and β represent the upper and lower bound of historical wind observations.

3 Experiments

Data and experimental setup. An open-source quadcopter simulator [4] is used to generate spatially and temporally varying winds. Three wind models are considered: (i) constant, (ii) sinusoidal, and (iii) Dryden gusts, a standard choice for UAM evaluation [7]. The wind field is updated at $f_w = 1$ kHz. As baselines, we implement standard kinematic wind estimators using Extended Kalman Filter (EKF) [8] and Unscented Kalman Filter (UKF) [12], which are widely used in commercial drone wind estimation and prove robust performance.

Table 1: Statistics of simulated wind scenarios (mean $\pm$ std). H: Hovering, C: Cruising.

Wind Model	**Magnitude (m/s)**	**Heading (rad)**	**Elevation (rad)**
Constant	10	1.05	0.26
Sinusoidal (H)	5.12$\pm$1.43	1.41$\pm$0.19	1.01$\pm$0.05
Sinusoidal (C)	4.67$\pm$1.43	1.72$\pm$0.79	0.77$\pm$0.05
Dryden (H)	6.77$\pm$2.13	1.06$\pm$0.45	0.13$\pm$0.44
Dryden (C)	8.24$\pm$2.02	1.42$\pm$0.34	0.75$\pm$0.24

Figure 2: Urban layout and flight path for the cruising scenario. The UAV departs from the Mudd building (red point), follows the red path through streets and courtyards to the riverside, circles the campus, and returns.

Both hovering and cruising cases are evaluated. The cruising case uses the Columbia University Morningside Heights campus (Figure 2), which induces diverse flows due to tall buildings, open spaces, and closeness to the Hudson River. The turbulence is modeled by the Dryden gust model

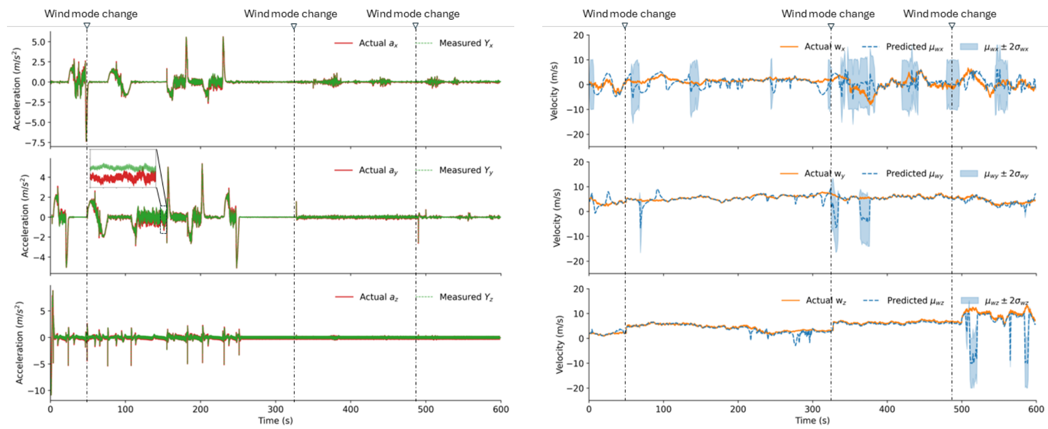

(a) Measured acceleration

(b) Predicted wind velocity

Figure 3: Cruising under time-varying Dryden gusts. (a) IMU measurements (affected by noise and bias); (b) predicted wind vector vs. ground truth, with uncertainty shown as shaded bands. Mode-change points are marked by dash-dotted lines.

with superposed step changes at street intersections (i.e., wind regime changes when moving from corridors to open spaces) to be more challenging and realistic.

Cruising trajectories are generated via the minimum-jerk method [6]. A PID cascade controls the x–y–z positions [4]. Maximum airspeed is 20 m/s; average cruising speed is 6 m/s. Each scenario runs for 10 minutes at $f_s = 100$ Hz. All simulations were executed on a Windows 10 machine with an NVIDIA A6000 GPU and an Intel Xeon 2.20 GHz CPU.

Main results. The root mean square error (RMSE) between the estimation and ground truth is used to evaluate the accuracy of the estimation, a measure widely used in previous studies (e.g., [14], [10]). Table 2 summarizes the results across various cases. It is shown that our method outperforms EKF and UKF in all scenarios. This demonstrates its ability to capture different wind modes during the flights stably, which is crucial for wind sensing in complex urban environments and supporting downstream tasks such as routing and path planning.

For more details, an example of wind estimation during cruising by our proposed method is shown below. During cruising, wind estimation becomes more challenging due to additional accelerations from UAV movement, introducing errors in data fusion, particularly under dynamic wind conditions. As shown in Figure 3(a), measured acceleration during cruising can be 4-5 times greater than in hovering, with larger deviations from actual states. This results in less accurate predictions in all three directions, as shown in Figure 3(b). However, accuracy improves significantly after a brief period of instability, where uncertainty is indicated by the blue shaded areas. This may be due to the UAV's rapid lifting and high-frequency, non-stationary motions, causing large measurement errors, which might be mitigated with warm-up. In fact, the method performs well later in the flight, with overall RMSE values of 2.42, 3.13, and 1.06 m/s (10%, 25%, 10% of speed range) in the x, y, and z directions, respectively.

4 Conclusion

We introduce a physics-informed probabilistic learner that estimates time-varying local horizontal and vertical wind components using only motion data from onboard IMU sensors with calibrated uncertainty. This framework is applicable for common quadrotor UAVs and does not need external sensors or complex dynamic models. Results across hovering and cruising show accurate recovery under varied wind regimes. This makes it suitable for predicting low-altitude wind variability. In addition, the framework quantifies the uncertainty of the model, which is critical for risk-sensitive applications such as fleet planning of UAVs and urban air mobility. The simplicity and effectiveness of the model suggest its potential for scalable deployment, enabling local-weather-aware UAV operations in cities.

Table 2: Comparison of RMSE (m/s) under different wind models for both hovering and cruising cases. The best-performing method for each scenario is highlighted in **bold**.

		Hovering			Cruising		
Wind Model	Method	x	y	z	x	y	z
Constant	Proposed	**0.17**	**0.34**	**0.11**	**0.80**	**1.09**	**1.49**
	EKF	0.22	0.41	0.14	0.96	1.30	1.80
	UKF	0.19	0.37	0.12	0.88	1.21	1.65
Sine	Proposed	**0.65**	**0.25**	**0.23**	**0.76**	**1.13**	**2.07**
	EKF	0.88	0.34	0.31	0.92	1.36	2.45
	UKF	0.74	0.29	0.26	0.84	1.25	2.28
Dryden	Proposed	**1.90**	**2.98**	**0.52**	**2.42**	**3.13**	**1.06**
	EKF	2.45	3.65	0.70	3.05	3.95	1.30
	UKF	2.12	3.25	0.60	2.74	3.50	1.18

References

[1] M. Chen, A. W. Smyth, M. G. Giometto, and M. Z. Li. Drone delivery routing with stochastic urban wind. In Proceedings of the 2023 IEEE 26th International Conference on Intelligent Transportation Systems (ITSC), pages 2260–2267. IEEE, 2023.

[2] P. Fontaine. Urban air mobility (uam) concept of operations v2.0. Technical report, FAA NextGEN, April 2023.

[3] M. Giometto, A. Christen, C. Meneveau, J. Fang, M. Krafczyk, and M. Parlange. Spatial characteristics of roughness sublayer mean flow and turbulence over a realistic urban surface. Boundary-Layer Meteorology, 160:425–452, 2016.

[4] GitHub. Quadcopter simulation and control. `https://github.com/bobzwik/Quadcopter_SimCon?tab=readme-ov-file`, 2019. Accessed: Sep. 14, 2024.

[5] T. Kane and D. Levinson. Dynamics, Theory and Applications. McGraw-Hill, 1985.

[6] K. J. Kyriakopoulos and G. N. Saridis. Minimum jerk path generation. In Proceedings of the IEEE International Conference on Robotics and Automation, pages 364–369, 1988.

[7] D. S. Nithya, G. Quaranta, V. Muscarello, and M. Liang. Review of wind flow modelling in urban environments to support the development of urban air mobility. Drones, 8(4):147, 2024.

[8] Ross T Palomaki, Nathan T Rose, Michael van den Bossche, Thomas J Sherman, and Stephan FJ De Wekker. Wind estimation in the lower atmosphere using multirotor aircraft. Journal of Atmospheric and Oceanic Technology, 34(5):1183–1191, 2017.

[9] C. Rose. Amazon's jeff bezos looks to the future. *CBS News*, 2013. Issue 1.

[10] T. Shimura, M. Inoue, H. Tsujimoto, K. Sasaki, and M. Iguchi. Estimation of wind vector profile using a hexarotor unmanned aerial vehicle and its application to meteorological observation up to 1000m above surface. Journal of Atmospheric and Oceanic Technology, 35(8):1621–1631, 2018.

[11] J. Stewart. Google tests drone deliveries in project wing trials. *BBC World Service Radio*, 2014.

[12] Andrea Tagliabue, Aleix Paris, Suhan Kim, Regan Kubicek, Sarah Bergbreiter, and Jonathan P How. Touch the wind: Simultaneous airflow, drag and interaction sensing on a multirotor. In 2020 IEEE/RSJ international conference on intelligent robots and systems (IROS), pages 1645–1652. IEEE, 2020.

[13] Qi Wang, Wei Wang, and Satoshi Suzuki. Uav trajectory tracking under wind disturbance based on novel antidisturbance sliding mode control. Aerospace Science and Technology, 149: 109138, 2024.

[14] A. Wenz and T. A. Johansen. Estimation of wind velocities and aerodynamic coefficients for uavs using standard autopilot sensors and a moving horizon estimator. In Proceedings of the International Conference on Unmanned Aircraft Systems (ICUAS), pages 1267–1276, 2017.

[15] Tamino Wetz, Norman Wildmann, and Frank Beyrich. Distributed wind measurements with multiple quadrotor unmanned aerial vehicles in the atmospheric boundary layer. Atmospheric Measurement Techniques, 14(5):3795–3814, 2021.

NeurIPS Paper Checklist

1. **Claims**

   Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

   Answer: [Yes]

   Justification: The main claims made in the abstract and introduction accurately reflect our paper's contributions and scope.

   Guidelines:

   - The answer NA means that the abstract and introduction do not include the claims made in the paper.
   - The abstract and/or introduction should clearly state the claims made, including the contributions made in the paper and important assumptions and limitations. A No or NA answer to this question will not be perceived well by the reviewers.
   - The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalize to other settings.
   - It is fine to include aspirational goals as motivation as long as it is clear that these goals are not attained by the paper.

2. **Limitations**

   Question: Does the paper discuss the limitations of the work performed by the authors?

   Answer: [Yes]

   Justification: Our limitations are discussed in the conclusions.

   Guidelines:

   - The answer NA means that the paper has no limitation while the answer No means that the paper has limitations, but those are not discussed in the paper.
   - The authors are encouraged to create a separate "Limitations" section in their paper.
   - The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
   - The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.
   - The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
   - The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
   - If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.
   - While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren't acknowledged in the paper. The authors should use their best judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

3. **Theory assumptions and proofs**

   Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

   Answer: [NA]

Justification: No proofs are needed.

Guidelines:

- The answer NA means that the paper does not include theoretical results.
- All the theorems, formulas, and proofs in the paper should be numbered and cross-referenced.
- All assumptions should be clearly stated or referenced in the statement of any theorems.
- The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
- Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
- Theorems and Lemmas that the proof relies upon should be properly referenced.

4. **Experimental result reproducibility**

   Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

   Answer: [Yes]

   Justification: We detail our training and evaluation process in the experimental setting section.

   Guidelines:

   - The answer NA means that the paper does not include experiments.
   - If the paper includes experiments, a No answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.
   - If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
   - Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general. releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
   - While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
     (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
     (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
     (c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).
     (d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

5. **Open access to data and code**

   Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

Answer: [No]

Justification: The data is confidential.

Guidelines:

- The answer NA means that paper does not include experiments requiring code.
- Please see the NeurIPS code and data submission guidelines (`https://nips.cc/public/guides/CodeSubmissionPolicy`) for more details.
- While we encourage the release of code and data, we understand that this might not be possible, so "No" is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
- The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (`https://nips.cc/public/guides/CodeSubmissionPolicy`) for more details.
- The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.
- The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
- At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).
- Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

6. **Experimental setting/details**

   Question: Does the paper specify all the training and test details (e.g., data splits, hyper-parameters, how they were chosen, type of optimizer, etc.) necessary to understand the results?

   Answer: [Yes]

   Justification: We detail our experimental setting in the Experiments Section.

   Guidelines:

   - The answer NA means that the paper does not include experiments.
   - The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
   - The full details can be provided either with the code, in appendix, or as supplemental material.

7. **Experiment statistical significance**

   Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

   Answer: [Yes]

   Justification: See Table 2.

   Guidelines:

   - The answer NA means that the paper does not include experiments.
   - The authors should answer "Yes" if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.
   - The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).
   - The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)
   - The assumptions made should be given (e.g., Normally distributed errors).
   - It should be clear whether the error bar is the standard deviation or the standard error of the mean.

- It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
- For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g. negative error rates).
- If error bars are reported in tables or plots, The authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

8. **Experiments compute resources**

   Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

   Answer: [Yes]

   Justification: We report the used computer resources in Experiments Section.

   Guidelines:

   - The answer NA means that the paper does not include experiments.
   - The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.
   - The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
   - The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

9. **Code of ethics**

   Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics `https://neurips.cc/public/EthicsGuidelines`?

   Answer: [Yes]

   Justification: The research conducted in the paper conforms, in every respect, with the NeurIPS Code of Ethics.

   Guidelines:

   - The answer NA means that the authors have not reviewed the NeurIPS Code of Ethics.
   - If the authors answer No, they should explain the special circumstances that require a deviation from the Code of Ethics.
   - The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

10. **Broader impacts**

    Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

    Answer: [Yes]

    Justification: We detail the potential impacts on urban air mobility in the conclusions.

    Guidelines:

    - The answer NA means that there is no societal impact of the work performed.
    - If the authors answer NA or No, they should explain why their work has no societal impact or why the paper does not address societal impact.
    - Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.

- The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.
- The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
- If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

11. **Safeguards**

    Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pretrained language models, image generators, or scraped datasets)?

    Answer: [NA]

    Justification: This paper does not have such risks.

    Guidelines:

    - The answer NA means that the paper poses no such risks.
    - Released models that have a high risk for misuse or dual-use should be released with necessary safeguards to allow for controlled use of the model, for example by requiring that users adhere to usage guidelines or restrictions to access the model or implementing safety filters.
    - Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
    - We recognize that providing effective safeguards is challenging, and many papers do not require this, but we encourage authors to take this into account and make a best faith effort.

12. **Licenses for existing assets**

    Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

    Answer: [Yes]

    Justification: The creators or original owners of assets (e.g., code, data, models), used in the paper, are properly credited. The license and terms of use are explicitly mentioned and properly respected.

    Guidelines:

    - The answer NA means that the paper does not use existing assets.
    - The authors should cite the original paper that produced the code package or dataset.
    - The authors should state which version of the asset is used and, if possible, include a URL.
    - The name of the license (e.g., CC-BY 4.0) should be included for each asset.
    - For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.
    - If assets are released, the license, copyright information, and terms of use in the package should be provided. For popular datasets, `paperswithcode.com/datasets` has curated licenses for some datasets. Their licensing guide can help determine the license of a dataset.

- For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.
- If this information is not available online, the authors are encouraged to reach out to the asset's creators.

13. **New assets**

    Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

    Answer: [NA]

    Justification: This paper does not have such issues.

    Guidelines:

    - The answer NA means that the paper does not release new assets.
    - Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
    - The paper should discuss whether and how consent was obtained from people whose asset is used.
    - At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

14. **Crowdsourcing and research with human subjects**

    Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

    Answer: [NA]

    Justification: This paper does not have such issues.

    Guidelines:

    - The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
    - Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
    - According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

15. **Institutional review board (IRB) approvals or equivalent for research with human subjects**

    Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

    Answer: [NA]

    Justification: This paper does not have such issues.

    Guidelines:

    - The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
    - Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.
    - We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
    - For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.

16. **Declaration of LLM usage**

    Question: Does the paper describe the usage of LLMs if it is an important, original, or non-standard component of the core methods in this research? Note that if the LLM is used only for writing, editing, or formatting purposes and does not impact the core methodology, scientific rigorousness, or originality of the research, declaration is not required.

    Answer: [NA]

    Justification: We do not involve LLMs as any important, original, or non-standard components in this paper.

    Guidelines:

    - The answer NA means that the core method development in this research does not involve LLMs as any important, original, or non-standard components.
    - Please refer to our LLM policy (`https://neurips.cc/Conferences/2025/LLM`) for what should or should not be described.