FLOW MATCHING-BASED ACTIVE LEARNING FOR RADIO MAP CONSTRUCTION WITH LOW-ALTITUDE UAVS

Hao Sun[1], Shicong Liu[1], Xianghao Yu[1], Ying Sun[2]

[1]City University of Hong Kong, Hong Kong, [2]Pennsylvania State University, PA, USA

ABSTRACT

The employment of unmanned aerial vehicles (UAVs) in the low-altitude economy necessitates precise and real-time radio maps for reliable communication and safe navigation. However, constructing such maps is hindered by the infeasibility of exhaustive measurements due to UAVs' limited flight endurance. To address this, we propose a novel active learning framework for low-altitude radio map construction based on limited measurements. First, a Plug-and-Play (PnP)-refined flow matching algorithm is introduced, which leverages flow matching as a powerful generative prior within a PnP scheme to reconstruct high-fidelity radio maps. Second, the generative nature of flow matching is exploited to quantify uncertainty by generating an ensemble of radio maps and computing the location-wise variance. The resulting uncertainty map guides a multi-objective candidate selection and then a trajectory is planned via utility-aware path search (UAPS), directing the UAV to the most informative locations while taking travel costs into account. Simulation results demonstrate that our method significantly outperforms the baselines, achieving more than a 70% reduction in normalized mean squared error (NMSE).

***Index Terms*—** Active learning, flow matching, Plug-and-Play, radio map, UAVs.

1. INTRODUCTION

The rapid development of the low-altitude economy, empowered by the widespread deployment of UAVs in logistics, surveillance, and wireless communications [1, 2], has created a strong need for accurate radio maps [3,4]. A radio map reveals the spatial distribution of channel characteristics such as received signal strength (RSS), which plays a vital role in ensuring reliable communication, efficient resource allocation, and safe UAV operations, especially in complex urban environments [5, 6].

However, building such radio maps is difficult. The large areas to be covered, combined with the limited flight time and battery capacity of UAVs, make exhaustive measurements infeasible. Thus, reconstructing a complete radio map calls for advanced sampling and construction methods.

Existing approaches can be categorized into two groups. Traditional methods leverage statistical models such as Gaussian processes (GPs) to interpolate sparse data. While GPs can estimate prediction uncertainty and provide guidance for sampling [7, 8], they fail to scale and adapt to complex scenarios. More recent studies utilize deep learning techniques, including autoencoders and U-Nets, to learn latent map structures from data [9–11]. Unlike GPs, these models do not naturally provide reliable uncertainty estimates, which are essential for guiding sampling in active learning. To address this, these models typically incorporate additional modules for uncertainty estimation, such as Bayesian neural networks or deep ensembles. However, these methods are computationally expensive and may not yield reliable uncertainty estimates.

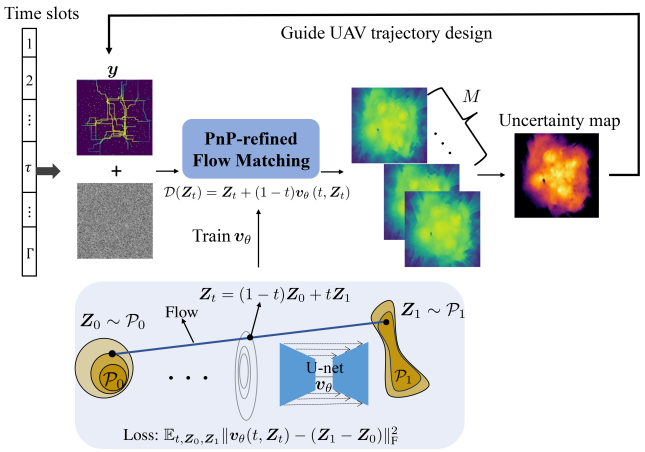

Fig. 1. Flow matching-based active learning process for radio map construction.

In recent years, the emergence of generative artificial intelligence (AI) offers an alternative, as its inherent stochasticity allows for the diverse outputs that naturally contain uncertainty, which serves as a direct and efficient measure of uncertainty. Yet, a unified framework that leverages generative diversity for both high-quality radio map reconstruction and uncertainty quantification remains to be explored.

In this paper, we develop a flow matching-based active learning framework for UAV-based radio map construction (cf. Fig. 1). We embed flow matching as a generative prior within a PnP scheme [12], and then introduce an adaptive refinement step that polishes the reconstruction in the final stages for higher fidelity. To obtain the uncertainty map, we exploit the generative nature of flow matching to generate multiple plausible radio maps and compute their location-wise variance. This uncertainty is then exploited to guide a probabilistic sampling strategy, which directs the UAV to informative yet cost-efficient measurement locations.

Our main contributions are summarized as follows:

- We propose a PnP-refined flow matching reconstruction scheme whose denoiser is defined by the flow matching velocity field and augmented with an adaptive inner-loop refinement.
- We quantify the uncertainty via a generative radio map ensembles obtained by repeatedly running the reconstruction process and computing the variance at each location.
- We design a probabilistic multi-objective selection rule that

balances informativeness and reachability, and plan a UAPS whose step cost embeds the normalized uncertainty.

2. PROBLEM STATEMENT

We consider a low-altitude urban environment where multiple base stations (BSs) are deployed, either at ground level or on building rooftops. The objective is to estimate the spatial distribution of the RSS over a low-altitude two-dimensional (2D) horizontal area, denoted by $\mathcal{S} \subset \mathbb{R}^2$, using a UAV. This area is discretized into a uniform grid of size $I \times J$, and the RSS value at each grid coordinate (i, j) is represented by the entry $\mathbf{M}_{ij}$ of a complete but unknown radio map matrix $\mathbf{M} \in \mathbb{R}^{I \times J}$.

A UAV is deployed to reconstruct $\mathbf{M}$ based on a limited number of measurements $\boldsymbol{y}$ collected along a trajectory across discrete time slots indexed by $\tau = 1, \ldots, \Gamma$. In each time slot τ, an ensemble of M candidate radio maps $\{\boldsymbol{Z}^{(\tau,m)}\}_{m=1}^{M}$ is generated by using a trained flow matching model based on all previously acquired measurements $\boldsymbol{y}^{(\tau)}$. From this ensemble of radio maps, an uncertainty map is derived to quantify the reliability of the estimated RSS values at each grid point. The UAV exploits this uncertainty map to determine its measurement trajectory in the next time slot $\tau + 1$.

3. PNP-REFINED FLOW MATCHING FOR RADIO MAP CONSTRUCTION

In this section, we study how to generate a candidate radio map $\boldsymbol{Z}$ in each time slot τ. For notation simplicity, the dependence on τ will be omitted throughout this section.

3.1. Radio Map Construction via a PnP Framework

The goal is to recover an unknown radio map $\boldsymbol{Z}$ from sparse RSS measurements $\boldsymbol{y} \in \mathbb{R}^Q$. This can be formulated as a maximum a posteriori (MAP) estimation

$$\arg\max_{\boldsymbol{Z}} \{\log f(\boldsymbol{y}|\boldsymbol{Z}) + \log p(\boldsymbol{Z})\}, \tag{1}$$

where $f(\boldsymbol{y}|\boldsymbol{Z})$ is the likelihood, and $p(\boldsymbol{Z})$ is the prior probability density function of the target radio map $\boldsymbol{Z}$.

Since the closed form of $p(\boldsymbol{Z})$ is unknown, we adopt the PnP framework, which approximates the prior $\log p(\boldsymbol{Z})$ with a learned denoiser $\mathcal{D}(\cdot)$. The reconstruction is then obtained through the PnP gradient-denoising iteration [12]

$$\begin{cases} \hat{\boldsymbol{Z}}^{(u)} = \boldsymbol{Z}^{(u-1)} - \gamma \nabla F(\boldsymbol{Z}^{(u-1)}; \boldsymbol{y}) \\ \boldsymbol{Z}^{(u)} = \mathcal{D}(\hat{\boldsymbol{Z}}^{(u)}) \end{cases}, \tag{2}$$

where the data-fidelity term $F(\boldsymbol{Z}; \boldsymbol{y}) \coloneqq -\log f(\boldsymbol{y}|\boldsymbol{Z})$. One example of F is $F(\boldsymbol{Z}; \boldsymbol{y}) = \|\mathcal{H}(\boldsymbol{Z}) - \boldsymbol{y}\|^2$, with $\mathcal{H} : \mathbb{R}^{I \times J} \to \mathbb{R}^Q$ denoting a degradation operator. The symbol ∇ denotes the gradient operator, and γ is the step size. The first step enforces measurement consistency, and the second step projects the estimation onto the manifold of ground truth radio maps.

The performance of PnP crucially depends on the denoiser $\mathcal{D}(\cdot)$, for which we employ flow matching as a powerful generative prior.

3.2. Flow Matching

Let $\mathcal{P}_1$ represent the target distribution over the space of radio maps $\mathbb{R}^{I \times J}$. Note that in (1), $p(\boldsymbol{Z})$ refers to the density function associated with $\mathcal{P}_1$. Assume that we only have access to a number of radio map samples from $\mathcal{P}_1$, but cannot access the distribution itself. The goal of flow matching is to build a neural network that can generate new samples from $\mathcal{P}_1$, given a training dataset of radio map samples. To address this, flow matching builds a probability flow $\{\mathcal{P}_t\}_{0 \le t \le 1}$, from a known source distribution $\mathcal{P}_0$, typically a Gaussian one [13], to the target distribution of radio map $\mathcal{P}_1$.

A straightforward choice of flow is the linear interpolation between a source sample $\boldsymbol{Z}_0 \sim \mathcal{P}_0$ and a target sample $\boldsymbol{Z}_1 \sim \mathcal{P}_1$

$$\boldsymbol{Z}_t = (1-t)\boldsymbol{Z}_0 + t\boldsymbol{Z}_1, \quad t \in [0, 1]. \tag{3}$$

The instantaneous velocity of this path is $u(t, \boldsymbol{Z}_t) = \mathrm{d}(\boldsymbol{Z}_t)/\mathrm{d}t = \boldsymbol{Z}_1 - \boldsymbol{Z}_0$. Flow matching trains a neural network $\boldsymbol{v}_\theta(t, \boldsymbol{Z}_t)$, parameterized by θ, to approximate $u(t, \boldsymbol{Z}_t)$, with the objective

$$\mathcal{L}(\theta) = \mathbb{E}_{t, \boldsymbol{Z}_0, \boldsymbol{Z}_1} \|\boldsymbol{v}_\theta(t, \boldsymbol{Z}_t) - (\boldsymbol{Z}_1 - \boldsymbol{Z}_0)\|_{\mathrm{F}}^2, \tag{4}$$

where t is uniformly distributed over $[0, 1]$, and $\|\cdot\|_{\mathrm{F}}$ denotes Frobenius norm.

The learned velocity field $\boldsymbol{v}_\theta(t, \boldsymbol{Z}_t)$ determines a time-dependent flow $\boldsymbol{\phi} : [0, 1] \times \mathbb{R}^{I \times J} \to \mathbb{R}^{I \times J}$

$$\frac{\mathrm{d}\boldsymbol{\phi}(t, \boldsymbol{Z}_0)}{\mathrm{d}t} = \boldsymbol{v}_\theta(t, \boldsymbol{\phi}(t, \boldsymbol{Z}_0)), \tag{5}$$

with initial condition $\boldsymbol{\phi}(0, \boldsymbol{Z}_0) = \boldsymbol{Z}_0$. By integrating (5) from $t = 0$ to $t = 1$, the initial noise map $\boldsymbol{Z}_0$ is smoothly transported along the learned flow to a final state $\boldsymbol{\phi}(1, \boldsymbol{Z}_0) = \boldsymbol{Z}_1$ that lies on the manifold of ground truth radio maps.

3.3. PnP-Refined Flow Matching Framework

To integrate the learned flow matching prior into the PnP framework in (2), we propose the PnP-refined flow matching framework detailed in **Algorithm 1**. The framework proceeds over K outer iterations. To align these discrete iteration steps with the continuous time variable of the flow, we define a discrete time sequence $t_k = k/K \in [0, 1]$.

After the velocity field $\boldsymbol{v}_\theta(t, \boldsymbol{Z}_t)$ is learned in (4), we define the denoiser $\mathcal{D}(\cdot)$ at each time step t_k as [14]

$$\mathcal{D}(\tilde{\boldsymbol{Z}}_{t_k}^{(u)}) \triangleq \tilde{\boldsymbol{Z}}_{t_k}^{(u)} + (1 - t_k)\boldsymbol{v}_\theta\left(t_k, \tilde{\boldsymbol{Z}}_{t_k}^{(u)}\right). \tag{6}$$

Each outer iteration k, corresponding to the time variable t_k, contains an inner refinement loop of U_k steps. Crucially, the number of refinements U_k increases with k, so that early iterations with fewer inner steps rapidly establish the coarse structure of the radio map, while later iterations with more inner steps gradually refine the estimate to recover fine details. This progressive refinement schedule is implemented through the following inner loop, which balances data fidelity and the learned generative prior.

Each of the U_k steps begins by steering the current estimate $\boldsymbol{Z}_{t_k}^{(u-1)}$ toward consistency with the measurements $\boldsymbol{y}$ via a gradient descent step on the data-fidelity term F as in (2). To mitigate artifacts and keep the trajectory aligned with the generative path, the intermediate result $\hat{\boldsymbol{Z}}_{t_k}^{(u)}$ from the data-consistency update is then regularized through Path Projection by interpolating it with the initial noise map $\boldsymbol{Z}_0$: $\tilde{\boldsymbol{Z}}_{t_k}^{(u)} \leftarrow t_k \hat{\boldsymbol{Z}}_{t_k}^{(u)} + (1 - t_k)\boldsymbol{Z}_0$. Finally, the Denoising step applies the learned denoiser $\mathcal{D}$ in (6), which leverages the model's velocity field to project the estimate onto the manifold of ground truth radio maps, effectively removing structural inconsistencies introduced by the data-consistency update.

Algorithm 1 PnP-Refined Flow Matching for Radio Map Reconstruction

Input: Observed RSS $\boldsymbol{y}$; an initial noise map $\boldsymbol{Z}_0 \sim \mathcal{P}_0$; number of outer steps K; step size γ; refinement schedule $\{U_k\}_{k=1}^{K}$
Output: Reconstructed map $\boldsymbol{Z}$
for $k = 1, \ldots, K$ **do** ▷ outer loop at $t_k = k/K$
  $t_k \leftarrow k/K$
  $\boldsymbol{Z}_{t_k}^{(0)} \leftarrow \boldsymbol{Z}_{t_{k-1}}$
  for $u = 1, 2, \ldots, U_k$ **do** ▷ adaptive inner refinement loop
    $\hat{\boldsymbol{Z}}_{t_k}^{(u)} \leftarrow \boldsymbol{Z}_{t_k}^{(u-1)} - \gamma \nabla F\Big(\boldsymbol{Z}_{t_k}^{(u-1)}; \boldsymbol{y}\Big)$ ▷ data-consistency update
    $\tilde{\boldsymbol{Z}}_{t_k}^{(u)} \leftarrow t_k\, \hat{\boldsymbol{Z}}_{t_k}^{(u)} + (1 - t_k)\, \boldsymbol{Z}_0$ ▷ path projection
    $\boldsymbol{Z}_{t_k}^{(u)} \leftarrow \mathcal{D}_{t_k}\big(\tilde{\boldsymbol{Z}}_{t_k}^{(u)}\big)$ ▷ apply the learned denoiser
  end for
  $\boldsymbol{Z}_{t_k} \leftarrow \boldsymbol{Z}_{t_k}^{(U_k)}$ ▷ result of outer step k
end for
return $\boldsymbol{Z} \leftarrow \boldsymbol{Z}_{t_K}$

4. ACTIVE LEARNING VIA UNCERTAINTY QUANTIFICATION

With a powerful PnP-refined flow matching framework for reconstruction, the next critical task is to intelligently guide the UAV's trajectory for the sampling in the next time slot $\tau + 1$.

4.1. Uncertainty Quantification with Generative Ensembles

A key feature of the PnP-refined framework is its generative capability, which enables the construction of an ensemble of plausible radio maps rather than a single reconstruction, all consistent with the measurements $\boldsymbol{y}^{(\tau)}$ collected up to time slot τ. Specifically, an ensemble of M reconstructions $\{\boldsymbol{Z}^{(\tau,m)}\}_{m=1}^{M}$ is generated by executing Algorithm 1 for M times, each initialized with a sample $\boldsymbol{Z}_0 \sim \mathcal{P}_0$.

Although each map is consistent with the acquired measurements, discrepancies arise in unmeasured regions where the model depends on its learned prior. The variation across ensemble members thus provides a principled measure of model uncertainty. The uncertainty map in time slot τ is then defined as the location-wise variance across the ensemble as

$$\boldsymbol{U}^{(\tau)} = \frac{1}{M} \sum_{m=1}^{M} \left(\boldsymbol{Z}^{(\tau,m)} - \frac{1}{M} \sum_{n=1}^{M} \boldsymbol{Z}^{(\tau,n)} \right)^2. \tag{7}$$

4.2. Active Sampling Strategy

Based on the uncertainty map $\boldsymbol{U}^{(\tau)}$, we first select several N candidate informative locations to be sampled. Then, we propose a UAPS approach to direct the UAV flying trajectory passing through all N candidates in the next time slot.

4.2.1. Candidate Locations Selection

We define $\mathcal{V}^{(\tau)}$ as the set of sampled locations, and $\mathcal{V}^{(\tau)}_{\text{unsampled}}$, as the set of unsampled locations in time slot τ.

A straightforward approach is to always select the location with the maximum uncertainty [8, 11]. However, such a greedy rule is prone to myopic behavior, as the UAVs may repeatedly fly to isolated high-variance points, thereby lengthening the flight trajectory, increasing energy consumption, and ultimately leading to inefficient trajectories.

To address this, we adopt a probabilistic multi-objective strategy. Instead of committing to a single maximizer, we sample a diverse set of candidate objectives that balance informativeness with reachability. Concretely, a set of N candidate locations $\mathcal{G}^{(\tau)} = \{\boldsymbol{g}_1^{(\tau)}, \ldots, \boldsymbol{g}_N^{(\tau)}\}$ is drawn without replacement from $\mathcal{V}^{(\tau)}_{\text{unsampled}}$, where $\mathbf{g}_n^{(\tau)} = (i_n^{(\tau)}, j_n^{(\tau)})$ is the grid coordinate. The selection criterion is given by the weight

$$w(\boldsymbol{g}_n^{(\tau)}) = \frac{\boldsymbol{U}^{(\tau)}_{i_n j_n}}{1 + \kappa d\left(\boldsymbol{g}^{(\tau)}_{\text{UAV}}, \boldsymbol{g}_n^{(\tau)}\right)}, \tag{8}$$

where $\boldsymbol{U}^{(\tau)}_{i_n j_n}$ is the uncertainty at location $\boldsymbol{g}_n^{(\tau)}$, $d(\cdot,\cdot)$ is the Manhattan distance, $\kappa \geq 0$ trades off informativeness $\boldsymbol{U}^{(\tau)}_{i_n j_n}$ versus $d(\cdot,\cdot)$, and $\boldsymbol{g}^{(\tau)}_{\text{UAV}}$ denotes the initial UAV position in time slot τ.

4.2.2. Visiting Order and Overall Trajectory Planning

In the following, we omit the superscript (τ) for notational simplicity. Given the N candidate locations in $\mathcal{G}$, we need to determine an efficient visiting order when designing trajectory. Let $\Pi(\mathcal{G})$ denote the set of all permutations of $\mathcal{G}$. For a given visiting order $\sigma \in \Pi(\mathcal{G})$, the total trajectory cost is

$$J(\sigma) = C\big(\boldsymbol{g}_{\text{UAV}}, \boldsymbol{g}_{\sigma(1)}\big) + \sum_{i=1}^{N-1} C\big(\boldsymbol{g}_{\sigma(i)}, \boldsymbol{g}_{\sigma(i+1)}\big). \tag{9}$$

The trajectory cost between two candidate locations is given by

$$C(\boldsymbol{g}_0, \boldsymbol{g}_L) = \min_{\pi:\, \boldsymbol{g}_0 \to \boldsymbol{g}_L} \sum_{l=0}^{L-1} c(\boldsymbol{g}_l, \boldsymbol{g}_{l+1}), \tag{10}$$

where $\boldsymbol{g}_0$ denotes the start candidate location, i.e., $\boldsymbol{g}_{\text{UAV}}$ or $\boldsymbol{g}_{\sigma(i)}$ in (9) and $\boldsymbol{g}_L$ denotes the end candidate location, i.e., $\boldsymbol{g}_{\sigma(1)}$ or $\boldsymbol{g}_{\sigma(i+1)}$ in (9), and $\pi : \boldsymbol{g}_0 \to \boldsymbol{g}_L$ denotes an arbitrary trajectory from $\boldsymbol{g}_0$ to $\boldsymbol{g}_L$. The per-step cost $c\big(\boldsymbol{g}_l, \boldsymbol{g}_{l+1}\big)$ is designed to reward the exploration of uncertain regions

$$c\big(\boldsymbol{g}_l, \boldsymbol{g}_{l+1}\big) = 1 - \beta \frac{\boldsymbol{U}_{i_{l+1} j_{l+1}} - \min(\boldsymbol{U})}{\max(\boldsymbol{U}) - \min(\boldsymbol{U})}, \tag{11}$$

where $\beta \in [0, 1]$ is a parameter controlling the incentive for exploration, and $\min(\boldsymbol{U})$ and $\max(\boldsymbol{U})$ denote the minimum and maximum values among all entries of $\boldsymbol{U}$, respectively. A higher uncertainty $\boldsymbol{U}_{i_{l+1} j_{l+1}}$ at the next step $\boldsymbol{g}_{l+1}$ leads to a lower step cost $c\big(\boldsymbol{g}_l, \boldsymbol{g}_{l+1}\big)$, encouraging the trajectory to pass through such informative locations.

The optimal order and trajectory can be obtained by $\sigma^* = \arg\min_{\sigma \in \Pi(\mathcal{G})} J(\sigma)$.

4.2.3. UAPS for Trajectory Planning Between Two Candidates

Directly solving for (10) by enumerating all possible trajectories is computationally intractable. Therefore, we propose an efficient search algorithm named UAPS to find the optimal trajectory $\pi : \boldsymbol{g}_0 \to \boldsymbol{g}_L$ from a starting point $\boldsymbol{g}_0$ to an objective $\boldsymbol{g}_L$.

Starting from the UAV position $\boldsymbol{g}_0$, UAPS repeatedly examines the neighboring grids. For each neighbor location $\boldsymbol{g}$, the evaluation function is defined as $\varphi(\boldsymbol{g}) = r(\boldsymbol{g}) + h(\boldsymbol{g})$. Here, $r(\boldsymbol{g})$ is the actual cost of the trajectory from the start location $\boldsymbol{g}_0$ to the candidate location $\boldsymbol{g}$, calculated by summing the per-step costs $c\big(\boldsymbol{g}_l, \boldsymbol{g}_{l+1}\big)$ in

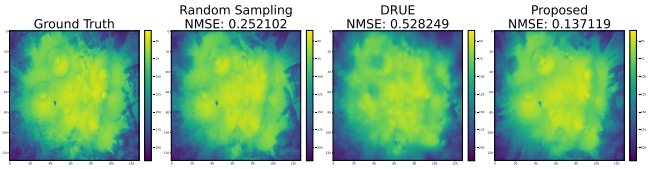

Fig. 2. Visual comparison of reconstruction radio maps.

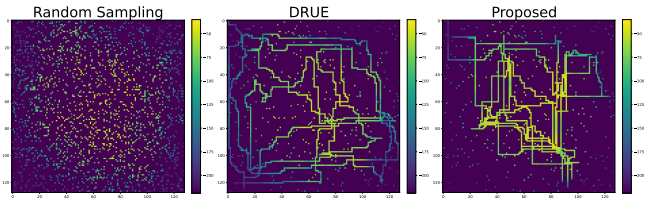

Fig. 3. Sampling patterns and UAV trajectories for Random Sampling, DRUE, and the proposed method.

(11) along the trajectory. The term $h(\boldsymbol{g})$ is a heuristic estimate of the remaining flying cost from the candidate location $\boldsymbol{g}$ to the objective $\boldsymbol{g}_L$, calculated as $h(\boldsymbol{g}) = (1-\beta)d(\boldsymbol{g}, \boldsymbol{g}_L)$. This formulation ensures that our search for an optimal trajectory correctly balances trajectory length and information gain in a computationally feasible manner.

At each location planning, UAPS expands the location with the smallest $\varphi(\boldsymbol{g})$, explores its feasible neighbors, and updates their costs. This process continues until the objective location $\boldsymbol{g}_L$ is reached. The resulting trajectory yields the travel cost (10), which is then employed in (9) to determine the optimal visiting order σ^*.

5. SIMULATION RESULTS

In this section, we present the simulation results to validate the effectiveness of our proposed flow matching-based active learning framework.

The ground truth radio maps were generated using the Sionna-RT ray-tracing simulator [15], which provides physically accurate radio channel predictions. We adopt the open-source "Etoile" 3D urban scenario, and randomly deploy 7 transmitters within this scenario. The ground truth radio map is then computed over a planar grid of 128×128 at a fixed altitude of 50 m, yielding an RSS matrix $\mathbf{M}$ with dimensions $I = J = 128$. This process is repeated to generate a dataset of 800 radio maps for training and 30 radio maps for testing. In the training stage of flow matching, the source distribution $\mathcal{P}_0$ is chosen as $\mathcal{N}(0, 1)$.

For demonstration, we assume that 2% of the grid locations are initially observed. The UAV then actively samples a total of $2,000$ additional locations, guided by the proposed active learning strategy. The degradation operator $\mathcal{H}$ models the UAV measurement process. $\mathcal{H}(\boldsymbol{Z})$ extracts the entries of the radio map corresponding to the locations that the UAV has visited. The PnP-refined flow matching algorithm is configured with $K = 50$ iterative steps and step size $\gamma = 2$. For steps $k < 46$, the refinement step U_k is set to 1, and increases to 10 in later iterations to enhance reconstruction precision. In the active learning phase, we set $\kappa = 0.001$ in (8) and exploration incentive $\beta = 0.9$ in (11). The number of candidate objectives per iteration is set to $N = 10$, and the ensemble size for uncertainty estimation $M = 5$.

The performance of the reconstruction is quantified using the NMSE. To ensure physical consistency, the NMSE is calculated in the linear power domain rather than the dB scale. The NMSE is then computed as $\text{NMSE} = \frac{\|10^{\mathbf{M}/10} - 10^{\boldsymbol{Z}/10}\|_{\text{F}}^2}{\|10^{\mathbf{M}/10}\|_{\text{F}}^2}$, where $\mathbf{M}$ and $\boldsymbol{Z}$ are the ground truth and reconstructed radio map in the linear power scale, respectively.

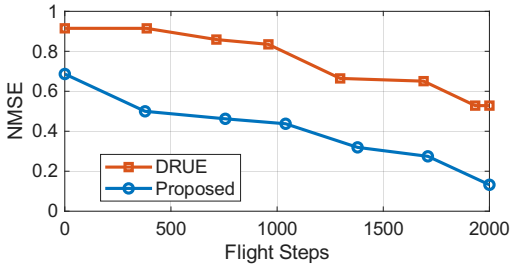

Fig. 4. Reconstruction NMSE versus the flight steps.

We benchmark our method against two sampling strategies, both constrained to the same overall sampling locations. Baseline 1: Random Sampling. In this baseline, $2,000$ additional sampling locations are selected uniformly at random from the entire grid. The reconstruction is then performed through the proposed PnP-refined flow matching based on all collected samples. This strategy does not involve an iterative, uncertainty-guided trajectory planning process. Baseline 2: Deep Radio Map and Uncertainty Estimator (DRUE) [11]. This method employs two autoencoders to estimate the radio map and its uncertainty separately, followed by a UAV trajectory designed in an uncertainty-aware manner.

The evaluation is conducted from three perspectives: qualitative reconstruction quality, sampling trajectory efficiency, and quantitative reconstruction error.

Fig. 2 provides a qualitative comparison of the radio map reconstruction after $2,000$ active samples. A visual inspection reveals that the radio map reconstructed by the proposed method more closely resembles the ground truth in both its detailed textures and overall spatial structure, qualitatively demonstrating its superior performance.

The superior reconstruction quality is a direct consequence of our intelligent sampling strategy, as illustrated in Fig. 3. While Random Sampling results in a scattered and inefficient coverage, and DRUE broadly explores the entire area without concentrating on truly informative regions, our method guides the UAV to precisely target the most informative regions.

The quantitative improvement is demonstrated in Fig. 4 which presents the reconstruction NMSE along the active learning process. It is clear that the NMSE of the proposed method consistently decreases at a faster rate than that of the DRUE baseline. Ultimately, our approach achieves an NMSE improvement of over 70% compared to the DRUE baseline, representing a significant performance gain.

6. CONCLUSION

This paper introduced an active learning framework for UAV-based radio map construction. A PnP-refined flow matching algorithm was introduced to enable high-fidelity map recovery. Radio map uncertainty was quantified via location-wise variance over generative ensembles, which subsequently guides a utility-aware trajectory planner to optimize UAV sampling trajectories. Simulation results verified that the proposed framework achieves superior reconstruction accuracy compared with baseline methods.

7. REFERENCES

[1] M. Wu, H. Wu, W. Lu, L. Guo, I. Lee, and A. Jamalipour, "Security-aware designs of multi-UAV deployment, task offloading and service placement in edge computing networks," *IEEE Trans. Mobile Comput.*, pp. 1–15, Oct. 2025.

[2] J. Chen, B. Li, H. Sun, S. Cui, and N. Pappas, "Predictive communications for low-altitude networks," Sep. 2025. [Online]. Available: https://arxiv.org/abs/2509.01705

[3] S. Timilsina, S. Shrestha, and X. Fu, "Quantized radio map estimation using tensor and deep generative models," *IEEE Trans. Signal Process.*, vol. 72, pp. 173–189, Nov. 2024.

[4] L. Xu, L. Cheng, J. Chen, W. Pu, and X. Fu, "Radio map estimation via latent-domain plug-and-play denoisers," in *Proc. IEEE Int. Conf. Acoustics, Speech, Signal Process.*, Suzhou, China, Apr. 2025, pp. 1–5.

[5] X. Mo, Y. Huang, and J. Xu, "Radio-map-based robust positioning optimization for UAV-enabled wireless power transfer," *IEEE Wireless Commun. Lett.*, vol. 9, no. 2, pp. 179–183, Oct. 2019.

[6] B. Li and J. Chen, "Radio map assisted approach for interference-aware predictive UAV communications," *IEEE Trans. Wireless Commun.*, Aug. 2024.

[7] Y. Chen, Q. Zhu, J. Wang, Z. Lin, Q. Wu, Y. Huang, and Q. Ye, "A novel online path planning method for UAV-based 3D spectrum mapping," in *Proc. IEEE Wireless Commun. Netw. Conf. (WCNC)*, Milan, Italy, Mar. 2025, pp. 01–06.

[8] K. D. Polyzos, A. Sadeghi, W. Ye, S. Sleder, K. Houssou, J. Calder, Z.-L. Zhang, and G. B. Giannakis, "Bayesian active learning for sample efficient 5G radio map reconstruction," *IEEE Trans. Wireless Commun.*, vol. 23, no. 12, pp. 19 382–19 396, Dec. 2024.

[9] Z. Chen, H. Wang, and D. Guo, "3-D radio map estimation based on active measurement trajectory selection," *IEEE Trans. Wireless Commun.*, vol. 14, no. 7, pp. 1884–1888, Jul. 2025.

[10] W. Lu, S. Gao, M. Wen, Y. Liang, L. Yang, C.-B. Chae, and H. V. Poor, "Bayesian-driven graph reasoning for active radio map construction," Aug. 2025. [Online]. Available: https://arxiv.org/abs/2508.09142

[11] R. Shrestha, D. Romero, and S. P. Chepuri, "Spectrum surveying: Active radio map estimation with autonomous UAVs," *IEEE Trans. Wireless Commun.*, vol. 22, no. 1, pp. 627–641, Jan. 2023.

[12] K. Zhang, Y. Li, W. Zuo, L. Zhang, L. Van Gool, and R. Timofte, "Plug-and-play image restoration with deep denoiser prior," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 44, no. 10, pp. 6360–6376, Oct. 2022.

[13] Y. Lipman, R. T. Q. Chen, H. Ben-Hamu, M. Nickel, and M. Le, "Flow matching for generative modeling," Feb. 2023. [Online]. Available: https://arxiv.org/abs/2210.02747

[14] S. Martin, A. Gagneux, P. Hagemann, and G. Steidl, "Pnp-flow: Plug-and-play image restoration with flow matching," in *Proc. Int. Conf. Learning Representations (ICLR)*, Singapore, Apr. 2025, poster track, ICLR 2025.

[15] F. A. Aoudia, J. Hoydis, M. Nimier-David, S. Cammerer, and A. Keller, "Sionna RT: Technical report," Apr. 2025. [Online]. Available: https://arxiv.org/abs/2504.21719